AD-A126 234 GENERALIZED MONITORING FACILITY USERS MANUAL CHANGE 4 (U) COMMAND AND CONTROL TECHNICAL CENTER WASHINGTON DC 01 JAN 83 CCTC-CSM-UM-246-82-CHG-4 1/3

UNCLASSIFIED F/G 9/2 NL

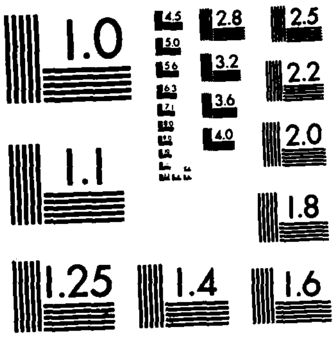

MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

(12)

DEFENSE COMMUNICATIONS AGENCY
COMMAND AND CONTROL TECHNICAL CENTER
WASHINGTON, D. C. 20301

AD A126234

IN REPLY REFER TO: C751

1 January 1983

TO: RECIPIENTS

SUBJECT: Change 4 to Computer System Manual-CSM UM 246-82, Generalized Monitoring Facility

1. This is Change 4 to Computer System Manual CSM UM 246-82, Generalized Monitoring Facility, dated 1 May 1982. Remove obsolete pages and destroy them in accordance with applicable security regulations and insert new pages as indicated below:

Remove Pages	Insert New Pages
Complete Table of Contents	Complete Table of Contents
xxviv	xxviv
1-1 and 1-2	1-1 and 1-2
Complete Section 2	Complete Section 2
Complete Section 3	Complete Section 3 (new)
Complete Section 4	Complete Section 4 (new)
5-5 through 5-27.3	5-5 through 5-27.3
5-27.20 and 5-28	5-27.20 and 5-28
5-43 and 5-43.1	5-43 and 5-44
5-43.2 and 5-44	5-45 and 5-46
5-44.1 and 5-44.2	5-46.1 and 5-46.2
5-45 and 5-46	5-53.22 through 5-70
5-53.22 through 5-68	6-3 through 6-16
6-3 through 6-16	6-19 and 6-19.1
6-19 and 6-19.1	6-21 and 6-22
6-21 and 6-22	6-29 through 6-40
6-29 through 6-40	7-45 through 7-58
7-44.3 through 7-56	9-29 and 9-29.1
9-29 and 9-29.1	Complete Section 11 (new)
Complete Section 11	14-15 through 14-24
14-15 through 14-24	14-31 and 14-32
14-31 and 14-32	14-33 and 14-34
	14-35 and 14-36

DTIC ELECTE MAR 28 1983 H

2. The effective date of these change pages is 1 January 1983.

83 03 28 032

CCTC Ltr, C751, Change 4 to Computer System Manual CSM UM 246-82, Generalized Monitoring Facility

3. When this change has been made, post an entry in the Record of Changes and file this letter before the title page.

FOR THE DIRECTOR:

115 Enclosures

ALAN GROSS
Colonel, USAF
Deputy Director
for Computer Services

CONTENTS

Section		Page

Section		Page

Section		Page
6.3.3.24	Report 24 - Delay Time in the System Scheduler	6-38
6.3.3.25	Report 25 - Delay Time Until Core Allocator . .	6-38
6.3.3.26	Report 26 through 31 - The Elapsed Time of a Busy State of the Processors	6-38
6.3.3.27	Report 32 - The Elapsed Time of a Busy State of Processors	6-38
6.3.3.28	Report 33 - Elapsed Time Between Allocator Calls in 1/100 of a Second	6-39
6.3.3.29	Report 34 - The I/O Time Charged per User Activity in Seconds	6-39
6.3.3.30	Report 35 - The CP Time Charged per User Activity in Seconds	6-39
6.3.3.31	Report 36 - The Number of Times a User Activity was Swapped	6-39
6.3.3.32	Report 37 - The Total Elapsed Time a User Activity was Swapped	6-39
6.3.3.33	Report 38 - The Number of Times a System Activity was Swapped	6-39
6.3.3.34	Report 39 - The Total Elapsed Time a System Activity was Swapped	6-39
6.3.3.35	Report 40 - Number of Extra Activities That Might Fit in Memory Using Compaction	6-39
6.3.3.36	Report 41 - Number of Extra Activities That Might Fit Memory Without Compaction	6-40
6.3.3.37	Report 42 - The Percent of Size-Time Product Used by a User Activity	6-40
6.3.3.38	Report 43 through 49 - The Length of Idle State in the Processors	6-40
6.3.3.39	Report 50 - The Original Allocation Time for User Memory in 1/10 Second	6-41
6.3.3.40	Report 51 - The Time Corrected Percent of Assigned Memory Used	6-41
6.3.4	Activity Resource Usage Report	6-41
6.3.5	SNUMB-IDENT Report	6-43
6.3.6	Memory Map Report	6-43
6.3.7	Demand List Report	6-48
6.3.8	Activity Abort Report	6-48
6.3.9	Jobs Out of Core Report	6-50
6.3.10	Excessive Resource Use Report	6-50
6.3.11	Allocation Status Report	6-53
6.3.12	Plots Reports	6-53
6.3.12.1	Plot 1 - Available Memory vs. Outstanding Demand in Core Allocator Queue vs. Outstanding Demand in Core Allocator + Peripheral Allocator Queue	6-56

Section		Page

Section		Page

Section		Page

Section		Page
9.5.10	Response Time Report	9-29
9.5.11	Error Messages	9-29.1
9.5.12	H6000-DN355 Reject Report	9-29.1
9.5.13	Abort Report	9-29.1
9.5.14	TS1 Initial Parameter Report	9-29.1
9.5.15	Mailbox Busy Report	9-29.1
9.6	Default Option Alteration	9-31
9.6.1	Timeframe Reduction Report (Action Code TIME)	9-31
9.6.2	Delta Timeframe Report (Action Code DELTA)	9-31
9.6.3	Histogram Report (Action Code HISTG)	9-33
9.6.4	Trace Dump Report (Action Code LIST or ALL)	9-33
9.6.5	Record Count Limitation (Action Code NREC)	9-33
9.6.6	Response Time Limit (Action Code RESP)	9-33
9.6.7	Think Time Limit (Action Code THINK)	9-33
9.6.8	Terminal Mailbox Dump (Action Code MAIL)	9-36
9.6.9	Terminal Busy Limit (Action Code BUSY)	9-36
9.6.10	W6.4/2H Data Reduction (Action Code RN)	9-36
9.6.11	Response Time Report Timeframe (Action Code RATE)	9-36
9.6.12	Response Time Report SNUMB (Action Code SNUMB)	9-36
9.6.13	Terminate Input Options (Action Code END)	9-36
9.6.14	Default Options	9-36
9.7	JCL	9-38
9.8	Multireel Processing	9-38
9.9	Tape Error Aborts	9-40
10.	DATANET-355 DATA REDUCTION PROGRAM (DDRP)	10-1
10.1	Introduction	10-1
10.2	Data Collection Methodology	10-1
10.3	Analytical Methodology	10-1
10.4	Data Reduction Methods	10-3
10.5	DDRP Output	10-4
10.5.1	DATANET-355 Tabular Reports	10-9
10.5.1.1	Average Response Time for all Users by DATANET (Report RESPD)	10-9
10.5.1.2	HSLA/Subchannel Transmission Report by DATANET (Report BUFF)	10-12
10.5.1.3	Average Response Time for Specifically Designated Line IDs (Report RESPL)	10-12
10.5.1.4	HSLA Subchannels Being Monitored (Report HSLA)	10-12
10.5.1.5	Card Images from I* of the GRTS-II Data Collection (Report CARD)	10-16
10.5.1.6	List of Active Line IDs	10-16
10.5.1.7	HSLA Threshold Report (Report THRESH)	10-16

Section		Page

Section		Page

ILLUSTRATIONS

Figure		Page

Figure		Page

Figure		Page
9-12	Response Time/User Think Time Limit Report	9-25
9-13	Terminal Session Report	9-26
9-14	High Terminal Usage Report	9-27
9-15	Opcode Count Report	9-28
9-16	Response Time Report	9-30
9-16.1	H6000-DN355 Reject Report	9-30.1
9-16.2	Abort Report	9-30.2
9-16.3	TS1 Initial Parameter Report	9-30.3
9-16.4	Mailbox Busy Report	9-30.4
9-17	Card Format - Input Option TIME	9-32
9-18	Histogram Reports, Input Option HISTG	9-35
9-19	Terminal Mailbox Dump, Input Option MAIL	9-37
10-1	Default Values For Plots	10-5
10-2	Average Response Time for All Uses by DATANET (Report RESPD)	10-10
10-3	HSLA/Subchannel Transmission Report (Report BUFF)	10-13
10-4	Average Response Time for Designated Lines (Report RESPL)	10-14
10-5	HSLA Subchannels Being Monitored (Report HSLA)	10-15
10-6	Card Images From I* FILE of the GMF Monitor (Report CARD)	10-17
10-7	List of Active Line IDs	10-18
10-8	HSLA Threshold Report (Report Thresh)	10-19
10-9	Number of Transactions Send to Host (Plot 1)	10-23
10-10	Number of 36 Bit Words Received From Host (Plot 4)	10-25
10-11	Percent Idle Time (Plot 5)	10-26
10-12	Percent Buffer Requests Denied (Plot 6)	10-27
10-13	Histogram of DATANET Response	10-29
10-14	Histogram Special Line ID Response	10-30
10-15	Plot Alteration (Action Code PLOT)	10-34
10-16	Set Timespan(s) For Processing (Action Code TIME)	10-36
10-17	Action Codes For Turning Reports Off	10-37
10-18	Action Codes For Turning Reports On	10-38
10-19	Turn Off Plots (Action Code NOPLOT)	10-41
10-20	Threshold Values (Action Code THRESH)	10-43
10-21	Turn Off Reports For Selected DATANETs (Action Code NORPT)	10-44
10-22	Debug Option (Action Code DEBUG)	10-46
10-23	Alter Histogram Parameters (Action Code CHANGE)	10-47
10-24	Unmatched/Duplicate Error Report (Action Code DUPLIC)	10-48
10-25	Reject Report (Action Code REJECT)	10-49
11-1	Sample Execution Report	11-4
11-2	CPU Title Page	11-7
11-3	CPU Queue Length Distribution	11-8

Figure		Page

Figure		Page

TABLES

SECTION 1. GENERAL

1.1 Purpose of the Users Manual

The Users Manual describes each of the programs in the Generalized Monitoring Facility (GMF), discusses input options required to run each program, and provides sample outputs generated by each program.

The Generalized Monitoring Facility is delivered on a FILSYS SAVE tape. A description of the software is presented in section 2. Installation procedures for the GMF Monitor and guidance on how to run GMF can be found in section 5 of this manual.

1.2 Project References

The Generalized Monitoring Facility was originally developed for the Government by Honeywell Information Systems. Since delivery of the completed software in 1975, the Computer Performance Evaluation Branch (C751) of the Command and Control Technical Center has extensively modified and rewritten the GMF system. Numerous software errors have been corrected and many new features have been added.

1.3 Terms and Abbreviations

The following abbreviations will be used throughout the document.

CAM	-	Communications Analysis Monitor
CM	-	Channel Monitor
CPU	-	Central Processing Unit
CPUM	-	CPU Monitor
FTS	-	File Transfer System (a program within the WWMCCS Network System)
GCOS	-	Generalized Comprehensive Operating System
GMC	-	Generalized Monitoring Collector
GMF	-	Generalized Monitoring Facility
GRTS	-	General Remote Terminal Supervisor
GRTM	-	GRTS Monitor
IDLEM	-	Idle Monitor

MSM - Mass Storage Monitor

MUM - Memory Utilization Monitor

NCP - Network Control Program (a program within the WWMCCS Network System)

RMC - Resource Monitor Collector

RMDRx - Resource Monitor Data Reduction Program 1 through 3

RMON - Resource Monitor

TM - Tape Monitor

TPEM - Transaction Processing Executive Monitor

TSS - Timesharing Subsystem

TSSM - Timesharing Subsystem Monitor

WWMCCS - World Wide Military Command and Control System

1.4 Security and Privacy

There are no classified data collected in the GMF, but there is one exception. If the Communications Analysis Monitor (CAM) is being run with the Specific Terminal Option, classified data may be collected.

1.5 Manual Format

Section 2 of this Manual provides an overview of the GMF system. A brief description of each program is included. Sections 3 through 12 and 15 describe the programs of the GMF system in detail. In the presentation the user will find the appropriate information to successfully operate each program. The format of the sections includes a discussion of each program in the input, processing, and output phases. Section 13 of the document describes the procedures that a user should follow if it is desired to create a new GMF monitor and section 14 provides detailed information as to how GMF should be used in conducting a complete system-wide evaluation.

SECTION 2. SYSTEM SUMMARY

2.1 System Application

This section explains the GMF system by logically grouping the programs into two subsystems, the General Monitor Collector (GMC) subsystem and the Resource Monitor Collector (RMC) subsystem. Overviews of the programs in both subsystems are provided. The purpose of the GMC subsystem is to provide a means of collecting detailed information on the operation of the operating system and on the flow of a job through the system. The purpose of the RMC subsystem is to provide a general view of the operation of the system. The RMC can be run on a daily basis to allow the continuous analysis of the system operation. When it appears a problem is occurring, the GMC can be run to further define and resolve the problem area.

2.2 System Operation

Both GMC and RMC monitor GCOS in real time and generate output tapes. These output tapes are then processed through a series of data reduction routines which produce histograms, graphs, and reports which allow an analyst to evaluate system performance. Both RMC and GMC have exclusive data reduction routines (i.e., a tape generated by RMC cannot be reduced by the GMC data reduction program). Differences between the GMC and RMC subsystems pertain to the methods used to collect measurement data.

The RMC is a sampling-based monitor. At 30-second intervals, the RMC enters execution, collects its data, and writes its records to the standard accounting file. Since it samples only at preselected intervals, it cannot present a complete history of what has occurred on the system. However, because the RMC is a sampling-based monitor, overhead is low. During the period of time when RMC is not in execution, it may be swapped out of memory.

The GMC is a trace-based monitor. The various monitors that are comprised in the GMC subsystem are called into execution by the occurrences of the events that are to be captured. The GMC can therefore be used to provide a complete and detailed history of system performances. Because all occurrences of a given event are retrieved, GMC overhead is higher than RMC. Unlike the RMC, GMC cannot be swapped or moved and must be locked into memory. Finally, the GMC requires its own dedicated tape drive for the writing of its generated records.

2.2.1 The Resource Monitor Collection (RMC) Subsystem. The RMC subsystem is composed of a data collector, RMCC, and three associated

data reduction routines. The RMC data collector is explained in section 3, and the RMC data reduction routines are discussed in section 4.

The RMCC is a sampling-based monitor. The RMCC will sample system queues and tables on a 30-second time interval. The data captured from these queues and tables are written to the system accounting file.

The output generated by the RMCC is input to a series of three data reduction routines that are responsible for generating all output reports. See figure 2-1 for the RMC system flowchart and figure 2-2 for the subroutines that comprise the RMC system.

2.2.2 The Generalized Monitor (GMC) Subsystem. The GMC subsystem is composed of a series of data collector programs and a related set of data reduction programs. Each data collector program consists of one or more subroutines, and each program is used to monitor a different area of system performance. The name of the program indicates the area of system performance measured. In addition, any combination of monitoring programs may be executed during a monitor session. A detailed description of the entire data collector facility is given in section 5. Figure 2-3 shows the interrelation of all programs within the GMC subsystem. Figure 2-4 shows the subroutines that comprise each data collector program and those trace types which need to be active for each subroutine.

The mechanism used by the GMC data collector for obtaining control from the operating system is that of the normal system trace. The trace records a history of the occurrence of one or more of 72 systems events, 65 of which are presently defined. This recording is done by the system executing a unique code set resident in the System Dispatcher Module (.MDISP). Execution of this code is common to all system trace events and provides the point at which the GMC obtains control. See section 5 for a detailed description of how GMC gains control from the system.

The executive routine of the GMC processes input cards for the data collection routines, determines which areas of the system are to be monitored, performs any necessary initialization, and controls all data buffering and tape writing. A detailed description of this routine is given in section 5. The associated GMC data reduction programs are described in sections 6 through 12 and section 15.

2.3 System Configuration

The GMF is designed to be run on a HIS 6000 computer system, running with WWMCCS GCOS release 6.4 or 7.2. These releases are equivalent to the HIS commercial 2H or 4JS (any level) GCOS releases. When GMC is

used on WWMCCS release 6.4, or commercial release 2H, the user must insure that the value for variable "SYS64" is changed from its current value of 0 to a value of 1. It should be noted that the CPU, TPE and TSS Monitors may be run only under a WWMCCS 7.2 release or commercial 4JS release. See subsection 2.6 for a complete description of all user requirements prior to using the GMC. The RMC can only be used with a WWMCCS 7.2 (commercial 4JS) GCOS release.

THIS PAGE LEFT INTENTIONALLY BLANK

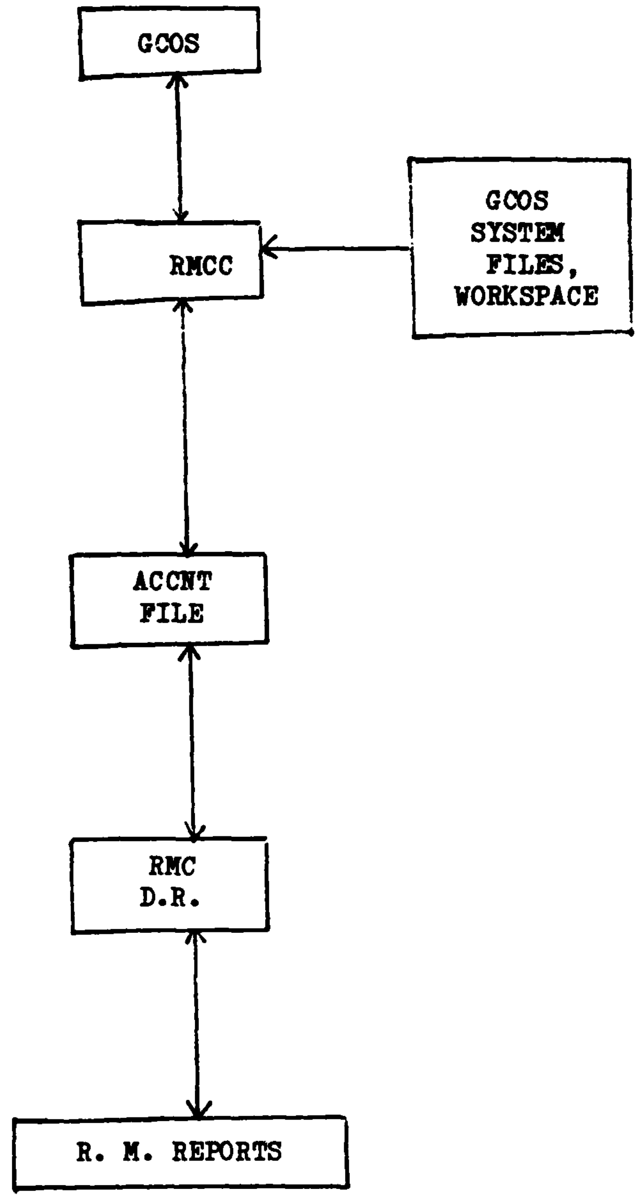

Figure 2-1. RMC Flowchart

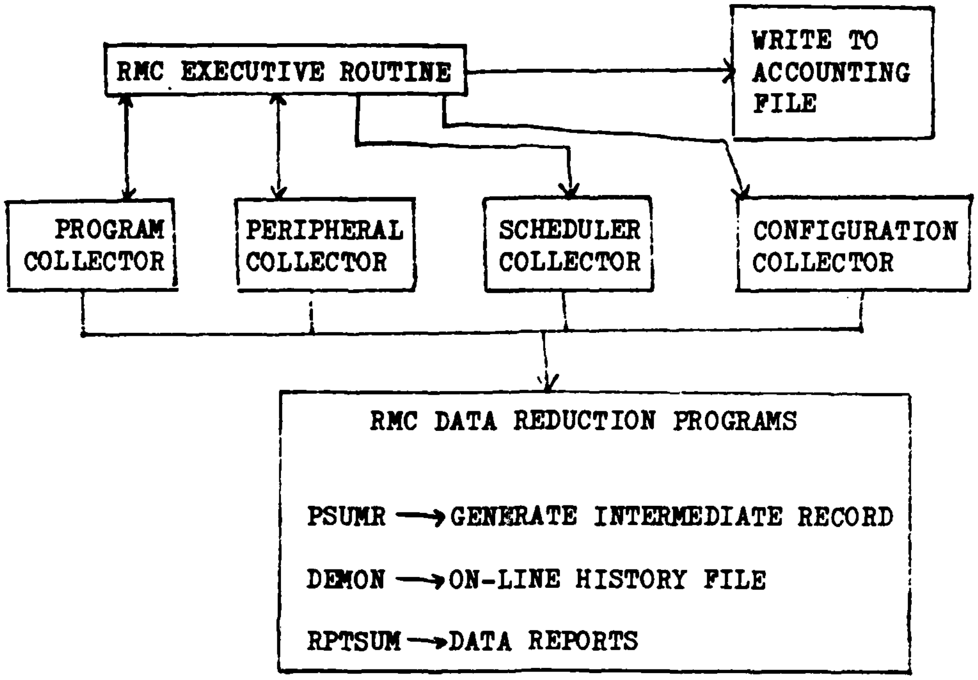

Figure 2-2. Programs in the RMC Subsystem

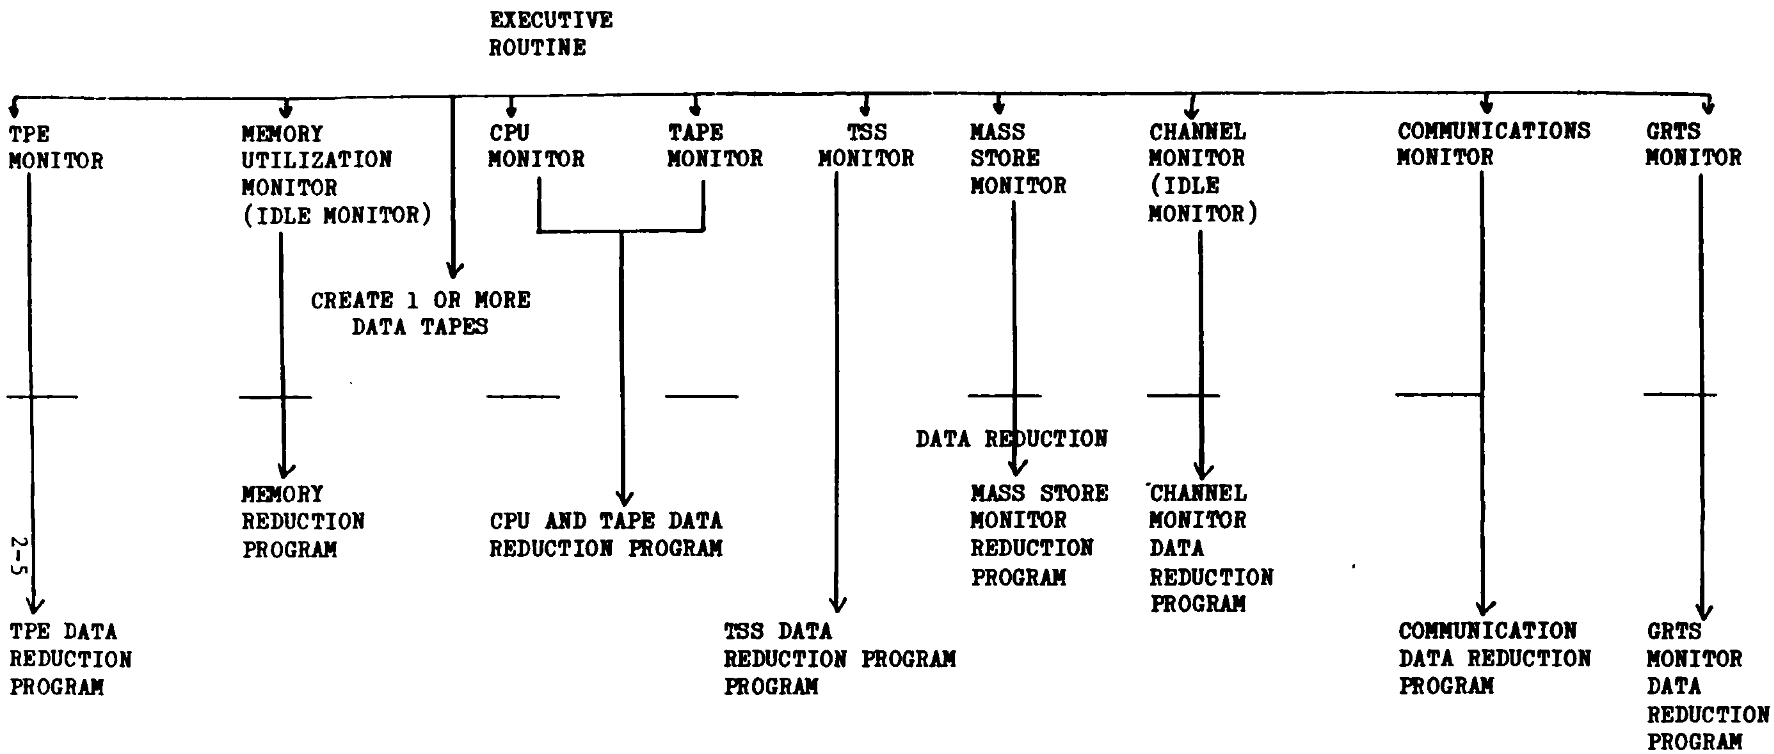

Figure 2-3. Programs in the GMC Subsystem

Data Collector Programs	Subroutines	Traces Captured (OCTAL) (NUMBER)
Memory Utilization Monitor	T10 T46	10,11,51 46
Idle Monitor	T21 TRCS	21 0,1,2,3,13,16,22, 37,65
Mass Store Monitor	T7	7,15,73*,76*,77*
Channel Monitor	T4,T7,T22	4,7,15,22
Tape Monitor	T50	50,51,52
CPU Monitor	T70	51,63*,70*
Communications Analysis Monitor	T14	14*,15
GRTS Monitor	T62	62*
TPE Monitor	T200	0,1,2,4,5,6,13, 42,51,65,74*
TSS Monitor	T100	74*

* - Nonstandard traces generated by the particular monitor.

Figure 2-4. Subroutines and Traces in GMC Data Collector Programs

2.4 System Organization

The GMF is composed of two data collectors, GMC and RMC, and associated data reduction programs. Sections 3 through 12 and section 15 describe these programs. Figure 2-1 gives a system flowchart for the RMC. Figure 5-1, from section 5, gives a system flowchart for the GMC.

2.5 Performance

The GMF monitors the performance of a system, aids in identifying the start of system performance problems, and aids in analyzing system performance problems. The RMC requires very little system resource usage and writes all its data to the system accounting file. The GMC is a much more detailed system with the associated higher overhead. The GMC is used mainly to determine the cause of system performance problems. The GMC requires 15 to 24 thousand words of memory and one tape drive while being run. Both systems require offline data reduction.

2.6 GMF Installation

2.6.1 Creation of GMF Files. The GMF software is contained on a single user save tape. The USERID on the tape is B29IDPXO. This USERID must be created with 5026 LLINKS of file space. A user restore can then be run. B29IDPXO is subdivided into several catalogs described below:

- GMFCOL - 520 LLINKS - This subcatalog contains all the data collection software for the GMC monitoring system. All files within this subcatalog are completely described in section 5.

- SOURCE - 2420 LLINKS - This subcatalog contains Time Sharing source files for all data reduction programs contained within the GMC system. Figure 2-5 is a breakdown of the individual files within this subcatalog. Sections 6-12 and 15 describe each program in detail.

- OBJECT - 1332 LLINKS - This subcatalog contains the object decks for all the data reduction programs contained within the GMC system. Figure 2-5 is a breakdown of the individual files within this subcatalog.

- JCL - 14 LLINKS - This subcatalog contains all the JCL required to run all the data reduction programs contained within the GMC system. Figure 2-6 is a breakdown of the individual files within this subcatalog.

- NEWRMON - 600 LLINKS - This subcatalog contains all the software required to collect and reduce the data for the RMON Monitoring

FILE NAME	FUNCTION	SOURCE SIZE (LL)	OBJECT SIZE (LL)
| MUM	MEMORY UTILIZATION MONITOR DATA REDUCTION PROGRAM. REFERENCED IN SECTION 6.	378	220
| MSM	MASS STORE MONITOR DATA REDUCTION PROGRAM. REFERENCED IN SECTION 7.	320	190
| CM	CHANNEL MONITOR DATA REDUCTION PROGRAM. REFERENCED IN SECTION 8.	255	182
| CAM	COMMUNICATION MONITOR DATA REDUCTION PROGRAM. REFERENCED IN SECTION 9.	190	100
| CPU-TAPE	CPU AND TAPE MONITOR DATA REDUCTION PROGRAMS. REFERENCED IN SECTION 11.	350	152
GRT	DATANET 355 MONITOR DATA REDUCTION PROGRAM. REFERENCED IN SECTION 10.	280	154
TPETG	TRANSACTION PROCESSING DATA REDUCTION PROGRAM. REFERENCED IN SECTION 12.	64	71
TPEALT	AN ALTER FILE FOR ADDING TPE TRACE CODE INTO THE TPE SUBSYSTEM (NO OBJECT FILE). REFERENCED IN SECTION 12.	14	
| TPEDUMP	A PROGRAM FOR OBTAINING A FORMATTED TRACE DUMP FROM A TPE/GMF DATA TAPE. REFERENCED IN SECTION 12.	40	26
| TSS	TIMESHARING MONITOR DATA REDUCTION PROGRAM. REFERENCED IN SECTION 15.	529	237

Figure 2-5. B29IDPXO/SOURCE and B29IDPXO/OBJECT Catalog Structure

FILE NAME	FUNCTION	SOURCE SIZE (LL)
MUM	JCL TO OBTAIN ALL MEMORY UTILIZATION MONITOR REPORTS	2
MSM	JCL TO OBTAIN MASS STORE MONITOR REPORTS	2
CM	JCL TO OBTAIN CHANNEL MONITOR REPORTS	2
CAM	JCL TO OBTAIN COMMUNI-CATION MONITOR REPORTS	2
GRT	JCL TO OBTAIN DN-355 MONITOR REPORTS	2
CPU-TAPE	JCL TO OBTAIN CPU AND TAPE MONITOR REPORTS	2
TPETG	JCL TO OBTAIN ALL REPORTS FROM THE TPE DATA REDUCTION PROGRAM	2
TSS	JCL TO OBTAIN ALL REPORTS FROM THE TSS DATA REDUCTION PROGRAM	2

Figure 2-6. B29IDPXO/JCL Catalog Structure

system. This subcatalog is further subdivided into JCL and SOURCE subcatalogs. The files within these subcatalogs are completely described in sections 3 and 4.

2.6.2 GMF Release Dependent Parameters. In order for the GMC to operate properly, it is necessary for GMC to locate certain instructions and/or words within several system programs. The user should insure that these locations are correct for the particular GCOS release, under which he is operating. Table 2-1 is a list of these dependent parameters identifying their the use and providing the approximate program source line numbers where the particular parameters are used. The list is provided for each GMC program that must checked by the user.

In order to use the GMC data reduction programs on a WW6.4 system, there is a special data card required in certain programs. This option applies to all data reduction programs except CPU-TAPE, TPETG and TSS. These programs are not designed for use under any release other than WW7.2 or commercial 4JS.

The GMC system is designed so that data collected on a WW6.4 system may be reduced under a WW6.4 system or a WW7.2 system. In addition, data collected under a WW7.2 system may likewise be reduced under a WW7.2 system, or a WW6.4 system. Whenever the data reduction programs for MUM, CM, CAM, or GRT are used on a WW6.4 system, a data card with an RN typed on it must be included in the input section of the JCL deck. It makes no difference under what release the data was collected. It is only a question of under what release the data is being reduced.

For the MSM data reduction program, there are two data cards required. The first data card always contains the letters RN. The second card is determined by the following table:

Data Collected	Data Reduced	Data Value
WW6.4	WW6.4	1
WW6.4	WW7.2	3
WW7.2	WW6.4	2
WW7.2	WW7.2	NO SPECIAL CARDS REQUIRED

Table 2-1. GMC Release Dependent Parameters
(Part 1 of 5)

Program	LINE #	Variable	Explanation
GMF.TOP	90	SYS64	Used to control conditional assembly of GMC set=1 for W6.4(2H) release set=0 for W7.2(4J) release
	10240	-	Code in this area searches for trace processing within the dispatcher. Trace code must be within 500 octal locations of the address specified by entry point 15 decimal of the dispatcher. This entry point should contain the address of location TRACE within the dispatcher, which is where the trace processing code is located. The code being searched for is an LDAQ;STAQ;TRAO,1.
	10760	-	Code in this area is used to make a correction to accounting processing, if the correction has not already been made via patches. The code is searched for within 500 octal locations of .MIOS entry point. The code searched for is SBLA TRREG+7,$;ARL 12; ADLA .CRTOD,7. The ARL is changed to an ARS.
CPU.PAT	240	-	Code in this area searches for an ASA .SALT,5 instruction in the dispatcher. Code must be within 300 octal locations after the address specified by entry point 20 decimal of the dispatcher. This entry point should contain the address of location DWAIT within the dispatcher.
	390	-	In a WW7.2 (4JS) system, code in this area searches for an STQ .QTOD,4 instruction. Search area is the same as described for line 240.
	540	-	If the dispatcher queue option of the CPU Monitor is activat., code in this area searches for an ORSA .STATE,4 instruction followed by an LDA .STATE,4 instruction. Code must be within 100 octal locations after the address specified by entry point 7 of the dispatcher. This entry point should contain the address of location DSPQH within the dispatcher.

Table 2-1. (Part 2 of 5)

Program	LINE #	Variable	Explanation
	680	-	In order to implant its special hooks, the CPU Monitor must modify the dispatcher and, therefore, it requires eight words of patch space. If the dispatcher queue option of the CPU Monitor is activated, then 12 words, instead of the normal 8, are required. This patch area must be within 200 octal locations in front of the address specified by entry point 15 decimal of the dispatcher. This entry point should contain the address of location TRACF within the dispatcher.
	1830	-	If sufficient patch space was not available in the standard patch area, the CPU Monitor will attempt to locate patch space in a specially defined user patch area. This search will take place only if bit 2 of word 0 within the dispatcher is set. This patch space should be within 200 octal locations after the address ILIST within the dispatcher. The address of ILIST is found in word 7 of the dispatcher.
	1580	-	Code in this area searches for an ARL 12 instruction, followed by an ASA .CROVH,7 instruction in the .MFALT module. The search for this code begins at the address specified by entry point 13 decimal of .MFALT and continues until the code is located. This entry point should contain the address of location BOOT within .MFALT.
CAM.INIT	930	-	Beginning at 1400 octal locations from the entry point of .MDNET and continuing for 5000 octal locations search for an ANA=0777777,DL (777777375207) instruction, followed by a CMPA (000000115210) instruction. This searches for number of special interrupts processed code (NSIP).

Program	LINE #	Variable	Explanation
	1190	-	Beginning at the CMPA location found above in .MDNET and continuing for 3000 octal locations search for an ANA=077 followed by a CMPA=077. When found, back up 30 octal words and look for an AOS instruction. This searches for the # of lines waiting mailbox code (R01XCT). If this code is not found, then the search is repeated looking for ANQ and CMPQ instructions instead. These instructions must be in an inhibited mode. This second search is required due to a redesign of the DNET module under commercial release 4JS3, edit level 4.
CAM.PAT	210	-	Code in this area searches for an LDQ M.LID,3 instruction, followed by an ANQ=0077777,DU instruction in module DNWW/DNET. The search for this code begins at the address specified by entry point -8 of DNWW/DNET and continues until the code is located. This entry point should contain the address of location NRQWT-DNET within DNWW/DNET.
	490	-	In order to implant its special hooks, the CAM must find 8 words of patch space. Even though the CAM is modifying module DNET, it will use the patch space available in the .MDISP module. This search for space is identical to that for CPU.PAT at lines 680 and 1830.
MSM.PAT	200	-	Code in this area searches the dispatcher for SSA cache code. If bit 4 of word 0 in the dispatcher is set, then cache is available. If this bit is not set, then no further searches are performed.
	250	-	Code in this area searches the dispatcher for the location of DBASE. The address at entry point -2 is obtained. This address points to the location ILIST in the dispatcher. At location ILIST, a series

Table 2-1. (Part 4 of 5)

Program	LINE #	Variable	Explanation
			of addresses are stored and MSM.PAT searches this list for the address DBASE.
	410	-	Code in this area searches for an AOS .CRTDL and an AOS .CRTBH instruction. This code needs to be within 300 octal locations after the address DBASE within the dispatcher.
	670	-	In order to implant its special hooks, the MSM Monitor must modify the dispatcher and, therefore, it requires 8 words of patch space. This search is identical to that for CPU.PAT at line 680 and line 1830.
GMF.MON	840	FMS1	Offset from entry point of .MFSIO which points to the word giving the absolute address of FMS catalog cache buffer. Used only in W7.2. Set to -13 decimal.
	850	FMS2	Offset from entry point of .MFSIO pointing to the word which gives the option selection for FMS catalog cache. Used only in W7.2. Set to -15 decimal.
MUM.T10	220	SYS64	See GMF.TOP
	920	FIFO	Address of the FIFO buffer within PALC. It is used to search the JCT table of PALC. This includes adding in a 110 octal offset for the loading of PALC in W6.4. There is no PALC offset in W7.2.
	5400	XPQ24	Location in CALC of the memory demand table. Set to octal 111.
	5410	SLVSNB	Offset in slave prefix area of job SNUMB. Set to octal 36.
	5420	MEMUSE	Offset in slave prefix area of loader memory use word. Set to octal 37.
	5430	IDENT	Offset in slave prefix area of job IDENT. Set to octal 66.

Table 2-1. (Part 5 of 5)

Program	LINE #	Variable	Explanation
CM.TO7A	190	IDENT	Offset in slave prefix area of job ident. Set to octal 66.
	210	SYS64	See GMF.TOP
	9600	-	This area of code searches .MFSIO in W7.2 to gather statistics for FMS catalog cache processing. Code should be checked to assure correct addresses are checked.
	11230	FFCCC	Address in PALC where the file code is stored during GEFSYE processing. Set to 6177 octal in W6.4 and 13143 octal in W7.2. This includes 110 octal for loading of PALC in W6.4. There is no offset for PALC in WW7.2.
	11240	SNUMBP	Address in PALC where the SNUMB is stored during GEFSYE processing. Set to 35012 octal in W6.4 and 2632 octal W7.2. This includes 110 octal for loading of PALC in W6.4. There is no offset for PALC in WW7.2.
	11250	ACT	Address in PALC where the activity number is stored during GEFSYE processing. Set to 33231 octal in W6.4 and 1051 octal in W7.2. This includes 110 octal for loading of PALC in W6.4. There is no offset for PALC in W7.2.

This page intentionally left blank

SECTION 3. RESOURCE MONITOR COLLECTOR

3.1 Introduction

The Resource Monitor Collector (RMC) is a privileged software package which periodically samples the system programs, queues, counters and tables of GCOS. The RMC consists of initialization routines, general purpose routines and four discrete data collector routines. The four data collectors include the following:

- o Program collector
- o Peripheral collector
- o Schedular collector
- o Configuration collector

Each of the data collectors generates a unique record format in one of the common buffer areas. A .CALL to .MSCF,2 is made to write each record to the system accounting file.

3.2 RMC Input Options and JCL

There are currently no RMC input options. The JCL needed to execute the RMC is shown in figure 3-1. Since the RMC runs in master mode, it requires a $PRIVITY card, meaning the operator must GRANT the job.

3.3 Processing

The RMC requires no special tapes or disk files. All data is written to the Statistical Collection File (SCF). The RMC will swap out of core if required and produce low system overhead. The next subsections discuss the collection routine of the RMC and give the record formats.

3.3.1 Collection Routine. The RMC collection routine is a GMAP program requiring PRIVITY. As currently released, the program uses SCF record number 608 for its data record.

The type 608 record consists of 1 of 4 subtypes. The record size depends on the system configuration. The system configuration record (subtype 4) is written once per hour. The RMC writes the other 3 type 608 records every 30 seconds. The RMC initializes its tables according to the system configuration. Extended memory instructions are NOPed if required, any excess memory is released and the ** file is released. The RMC then begins normal processing writing 3 SCF records every 30 seconds.

The following procedure applies to the WWMCCS software releases. The user should ensure that the 600 class of SCF records has been turned on for his

```
$   IDENT
$   USERID
$   PROGRAM COLL
$   PRMFL **,R,R,B29IDPXO/NEWRMON/DDYN
$   LIMITS 999,10K,,500
$   PRIVITY
$   ENDJOB
```

Figure 3-1. RMC JCL

system. He must check two different places. First, the $SCFBUF card in the boot deck must contain at least a C6 to indicate collection of 600-level SCF records, or else the specific record type 608 should appear on the card.

The user must next ensure that the system MASK cards are defined correctly. The 600-class SCF records must be turned on in INIT by use of a MASK card in the boot deck. The location to be changed with the MASK card is RMAX+6 in INIT. The location should be changed with the following MASK card:

COL 1	8	13	73
OCTAL LOCATION OF RMAX+6	MASK	000012000000	.MINIT

This card allows 10-type 600 SCF records to be collected (600-609). These same procedures should be followed by the user if he also changes the RMON SCF record number. Commercial 4Jx sites should check the System Startup Manual for procedures on defining new SCF record types. The user should then enter the command SCFRST 608 at the system console. The system should respond with: 608/AC to indicate that the records are being collected.

When used at a commercial site, the user must first execute the file B29IDPXO/NEWRMON/JCL/CUPJCL. Upon the successful completion of this step, the commercial user should then follow all standard operating procedures.

3.3.1.1 Program Data - Subtype 1. This data type contains overhead and idle times for all configured processors, size of memory, number of processors and number of TSS users. Each job in the system is examined and its status is saved. The subtype 1 record format is:

Word	Bits	Content	Source
0	0-17	Size	
	18-35	Record Type	
1-9	0-35	Standard SCF header	
10	0-29	Record sequence number	
	30-35	Subtype (=1)	
11	0-17	Current available memory (.CRPMU)	
	18-35	Number TSS users (.TCNOU)	
12	0-17	Unused	
	18-35	Number remotes connected (.CRCGT)	
13	0-35	System processor overhead (sum of all entries in .CROVH table)	
14	0-35	Processor idle time (sum of all entries in .CRIDT table)	
15	0-29	SNUMB (.CRSNB)	
	30-35	Unused	
16	0-17	Status (CRSN1)	
	18-35	Memory size if in memory (.SMSC)	

17	0-35	Accumulated processor time = -1 = not in memory (.SPRT)

Words 15-17 are repeated for all jobs in the system.

3.3.1.2 Peripheral Data - Subtype 2. This data type consists of channel use time for each channel, device status (released, assigned, dedicated, permanent, removable), and capacity of disk packs. The subtype 2 record format is:

Word	Bits	Content	Source
0	0-17	Size	
	18-35	Record type	
1-9	0-35	Standard SCF header	
10	0-29	Record sequence number	
	30-35	Record subtype (=2)	
11	0-35	Number available LLINKS for PRM devices (AST)	
12	0-35	Number available LLINKS for RMVB devices	
13	0-17	Number of two-word entries which follow	
	18-35	Not used	
14	0-5	Device type code (.CRSCT)	
	6-11	Number of removable devices on this channel (disk devices)	
	12-17	Number free	
	18-23	Number allocated	
	24-29	Number dedicated	
	30-35	Number released	
15	0-35	Channel use time (channel module + 8)	

Words 14-15 are repeated for each channel.

3.3.1.3 Scheduler Data - Subtype 3. This data type consists of system scheduler data. The subtype 3 record format is:

Word	Bits	Content	Source
0	0-17	Size	
	18-35	Record type	
1-9	0-35	Standard SCF header	
10	0-29	Record sequence number	
	30-35	Record subtype (=3)	
11	0-17	Number jobs allowed in class (CLSCAT)	
	18-35	Number jobs catalogued	
12	0	=1 = JEND on class (CLSJOB)	
	1-17	= Max number concurrent allowed	
	18-35	Number currently running	

Words 11-12 are repeated for each scheduler class

11+2n+1	0-35	= -1 fence

3.3.1.4 Configuration Data - Subtype 4. This data type consists of system configuration data. The subtype 4 record format is:

Word	Bits	Content	Source
0	0-17	Size	
	18-35	Record type	
1-9	0-35	Standard SCF header	
10	0-29	Record sequence number	
	30-35	Record subtype (=4)	
11	0-2	Number of IOMS	
	3-5	Number of processors	
	6-23	Configured memory	
	24-35	Hard core in KWDS	
12	0-8	Number tape channels	
	9-17	Number tape drives	
	18-26	Number printer channels	
	27-35	Number printers	
13	0-35	Number disk channels	
14	0-35	Number disk information entries	
15	0-5	Device type code	
	6	=1 = removable	
	7-17	Allocation unit	
	18-35	Maximum LLINKS for device	
16	0-17	Number defective LLINKS	
	18-35	Available table units	

Words 15-16 are repeated for each disk device configured.

3.4 Outputs

The RMC puts all of its output to the SCF.

3.5 Catalog Description

The RMC and data reduction files are located under the catalog B29IDPXO/NEWRMON. The resource monitor uses FILEDIT to maintain the source and object files. The programs are executed from a ** file maintained by SYSEDIT.

3.5.1 Files

3.5.1.1 FILEDIT Source File. The FILEDIT source file for all programs is:

B29IDPXO/NEWRMON/DSRC

3.5.1.2 FILEDIT Object File. The FILEDIT object file for all programs is:

B29IDPXO/NEWRMON/DOBJ

3.5.1.3 SYSEDIT Loadable File. The SYSEDIT loadable file (Q*) for all programs is:

B29IDPXO/NEWRMON/DDYN

3.5.1.4 JCL. The JCL to update and execute all programs is found under the following catalog:

B29IDPXO/NEWRMON/JCL

3.5.1.4.1 DUPJCL. File DUPJCL contains the JCL to update the source, object, and system loadable files.

3.5.1.4.2 DALT. File DALT will contain the FILEDIT directives when an update is desired. The user must place the directives on DALT prior to running DUPJCL.

3.5.1.4.3 COLL. File COLL contains the JCL for executing the collector.

3.5.1.4.4 PSUMR. File PSUMR contains the JCL for executing the intermediate record create program (PSUMR) and the daily monitor (DEMON).

3.5.1.4.5 RPTSUM. File RPTSUM contains the JCL for executing the report program (RPTSUM).

3.5.1.4.6 CUPJCL. File CUPJCL contains the JCL to update the source, object and system loadable files for commercial execution of the RMON system. This file must be processed prior to commercial use of this monitor.

3.5.2 Update Procedures. When changes are required to any of the resource monitor programs (besides those required to make RMON compatible to a commercial 4JS release), the FILEDIT directives are saved on B29IDPXO/NEWRMON/JCL/DALT. The user then runs the JCL found under file DUPJCL.

Sample TSS session to update the collector (after logon):

```
User:    Card N
System:  *
User:    10 $:MODIFY:source,object,COLL
         20 $:GMAP:::COLL
         30 $:UPDATE:LIST
         40 $:ALTER:10,10
            EOM
System:  *
User:    Resave B29IDPXO/NEWRMON/JCL/DALT
```

System: Data saved - DALT
User: OLD B29IDPXO/NEWRMON/JCL/DUPJCL
System: *
User: Run
System: Card format, disposition
User: S
System: Tab character, settings
User: N
System: SNUMB = XXXXT
User: Bye.

THIS PAGE LEFT INTENTIONALLY BLANK

SECTION 4. RESOURCE MONITOR DATA REDUCTION

The Resource Monitor Data Reduction (RMDR) is composed of three discrete programs. The first program (PSUMR) processes the SCF records producing an intermediate record. The second program is a daily monitor (DEMON) which maintains an on-line history file. The third program (RPTSUM) produces reports and prints the required plots. Figure 4-1 is an overview of the RMDR.

4.1 PSUMR

PSUMR selects the collector data records from an SCF input file. The records are sorted by system ID, date, time, and record subtype. The records are then processed in pairs to produce an intermediate record.

4.1.1 PSUMR Inputs. PSUMR requires one input file containing SCF records. Two optional inputs are the intermediate record - old master, and a parameter file.

4.1.1.1 SCF Records (File IN). Contains resource monitor collector (RMC) SCF records.

4.1.1.2 Intermediate Record - Old Master (File OM). This optional input will be copied to the output file before new records are added from the current jobs processing.

4.1.1.3 Parameter File (File PF). PSUMR will process the following parameter language statements:

CONFIG
SUMMARY
SNUMB

The format and use of the parameter language is found in section 4.4.

4.1.2 PSUMR Output. Three outputs are available from PSUMR: (1) listing of the parameter file if present; (2) intermediate record - new master file; (3) intermediate record - current file.

4.1.2.1 Parameter File Listing (File RP). This is a listing of the contents of file PF. Any of the statements processed which were invalid will be flagged. Invalid contents will not halt the job execution, but may affect how the data is summarized in the intermediate record.

4.1.2.2 Intermediate Record New - Master File (File NM). This file contains the contents of the old master file, if present, plus the records added from the current execution of the program.

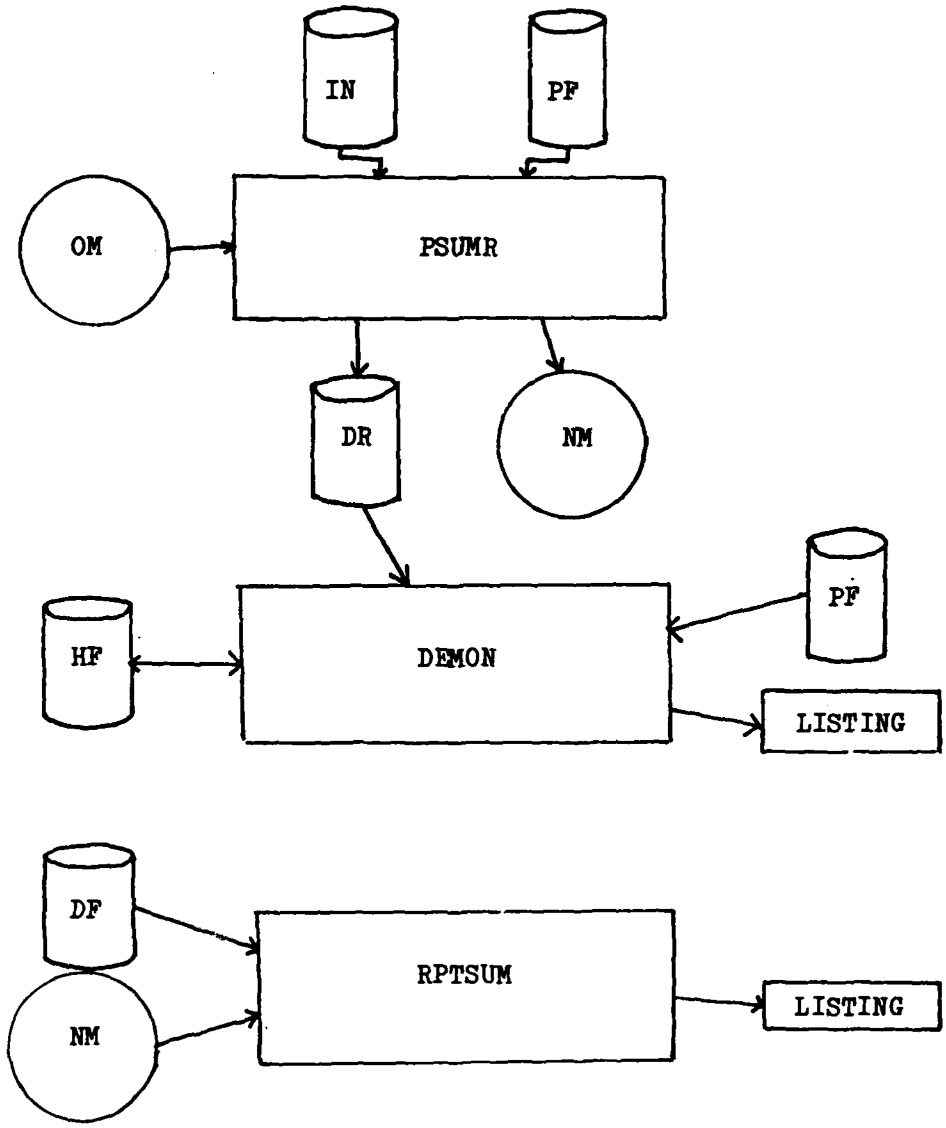

IN = SCF accounting records
PF = optional parameter file
OM = old master intermediate record
NM = new master intermediate record
DR = optional intermediate output for DEMON
HF = daily monitor history file

Figure 4-1. RMDR Overview

4.1.2.3 Intermediate Record - Current File (File DR). This file contains the intermediate records created during the current execution of the program. This file is used to pass the intermediate records to the daily monitor (DEMON), when present, in a subsequent activity.

4.1.3 PSUMR Deck Setup. The following control cards are required to execute PSUMR:

```
$      IDENT        ACCOUNTING INFORMATION
$      USERID       USERID$PASSWORD/SCC
$      PROGRAM      PSUMR,Dump
$      LIMITS       20,20K,,5K
$      PRMFL        **,R,R,B29IDPXO/NEWRMON/DDYN
$      DATA         PF                optional parameters
$      TAPE         OM                optional old master intermediate
                                      record
$      TAPE         NM                new master intermediate record
$      FILE         DR                current file output
$      SYSOUT       RP                parameter file listing
$      TAPE         IN                SCF input file
$      FILE         S1,S1R.50R        sort files
$      FILE         S2,S2R,50R
$      FILE         S3,S3R,50R
$      FILE         S4,S4R,50R
$      ENDJOB
```

The sort files' size should be increased or decreased depending upon the input volume. Tape files may be replaced by permanent or temporary disk files. File DR may be null if the daily monitor does not follow in a subsequent activity.

4.2 DEMON

The daily monitor accepts the intermediate records created by PSUMR. A history file will be initialized (if current null) or updated from the data in the intermediate record. A matrix structure is built to summarize and store the data. The matrix contains 96 entries corresponding to the time period to be summed together. The entry size is variable depending upon the graphs to be produced. For DEMON, the default is all graphs. The history file contains a copy of this matrix structure. For an update run, DEMON copies the history file into core. A second matrix is built for the current day's data. Both matrices are then updated from the input. The history matrix is written to disk. Finally, the current input matrix is compared to the history matrix. If any data item is not within 10% of the history matrix value, the program generates a complete set of graphs to indicate a significant change from the history file value.

4.2.1 DEMON Inputs. DEMON has one required input and two optional inputs.

4.2.1.1 History File (File HF). The history file is required for DEMON execution. The history file is a random file which contains an image of the in-core matrix used for summarizing and comparing data.

4.2.1.2 Parameter File (File PF). The parameter file is optional. However, if PSUMR was supplied with a SUMMARY statement, the DEMON should be supplied the same summary statement for desired results. DEMON processes only a summary statement. The format and use of the parameter language statements is found in section 4.4.

4.2.1.3 Intermediate Records (File IN). This is an optional input. If present, the intermediate record data will be summarized, as described above, according to parameter setting marked on the intermediate record. The history file will be initialized or updated as it is appropriate. If file IN is not present, the history file data will be used to generate a set of graphs. This feature allows the user to get a current historical "picture" at any time.

4.2.2 DEMON Outputs. DEMON outputs are the updated history file and a listing file.

4.2.2.1 Updated History File (File HF). Refer to section 4.2.1.3.

4.2.2.2 Listing File (File RP). The output listing will contain a listing of the parameter file, if any, and any generated graphs. The graphs produced are listed and described in section 4.5..

4.2.3 DEMON Deck Setup. The following control cards are required to execute DEMON.

```
$     IDENT         ACCOUNTING INFORMATION
$     USERID        USERID$PASSWORD
$     PROGRAM       DEMON,DUMP
$     LIMITS        10,24K,,10K
$     PRMFL         **,R,R,B29IDPXO/NEWRMON/DDYN
$     DATA          PF       optional parameter file
$     TAPE          IN       optional intermediate file
$     PRMFL         HF,W,R, (history cat/file)
$     SYSOUT        RP
$     ENDJOB
```

File IN may be a permanent or temporary disk file. File HF is 24 random LLINKS.

4.3 RPTSUM

PSUMR has the ability to maintain a master file of intermediate records. RPTSUM will process this file to produce any number of reports. Each report

consists of one to eight graphs covering arbitrary time frames. The user specifies, in a parameter file, the reports required.

4.3.1 RPTSUM Inputs. RPTSUM has one required input and one optional input.

4.3.1.1 Intermediate Master Record (File IF). Typically, this required file would contain records covering an extended time (weeks or months). RPTSUM will identify the records which fall within the time limits for each report. The file is processed only once per execution, regardless of the number of reports required.

4.3.1.2 Parameter File (File PF). RPTSUM produces reports based on statements from an optional parameter file. If a parameter file is not supplied, all data on the IF file will be processed. The format and use of the parameter language is found in section 4.4.

4.3.2 RPTSUM Outputs. RPTSUM has two output listing files.

4.3.2.1 Parameter File Listing (File RP). This is a listing of the parameters found on file PF, if any.

4.3.2.2 Report Listing (File GP). This listing has the printed reports requested, including the banner page, configuration overview and plots.

4.3.3 RPTSUM Deck Setup. The following control cards are required to execute RPTSUM:

```
$     IDENT       ACCOUNTING INFORMATION
$     USERID      USERID$PASSWORD
$     PROGRAM     RPTSUM,DUMP
$     PRMFL       **,R,R,B29IDPXO/NEWRMON/DDYN
$     TAPE9       IF      intermediate record input
$     DATA        PF      optional parameter file
$     SYSOUT      RP      parameter file listing
$     SYSOUT      GP      report listing
$     LIMITS      20,24K,,10K
```

The memory requirements depend on the number of reports and number of graphs requested for each report.

4.4 Parameter Language

The parameter language developed for PSUMR, DEMON and RPTSUM consists of free format statements. Each statement has a statement identifier, followed by one or more keyword phrases. A keyword phrase consists of subkeyword and/or user-supplied parameters.

In the following discussion, a set of closed parentheses is used to identify user-supplied parameters. The parentheses are not supplied by the user.

Closed brackets indicate optional parameters. Number signs surrounding a parameter imply that a parameter is required.

4.4.1 Statement Identifiers. The language contains five statement identifiers.

FORMAT	PURPOSE	PROCESSED BY
REPORT	specify report options	RPTSUM
CONFIG	specify standard configuration	PSUMR
SUMMARY	specify default include/exclude parameters	PSUMR DEMON
SNUMB	specify SNUMB to include in nonstandard categories	PSUMR
COMPARE	specify reports to compare	RPTSUM

4.4.2 Report Statement Keywords and Parameters. The keywords for the report statement are NAME, SYSTEM, PERIOD, TIME, INCLUDE and EXCLUDE.

4.4.2.1 NAME. This is used to supply an identifier for the report.

Format

NAME (report ID)

Comments

- maximum of six characters
- default = DEFALT

4.4.2.2 SYSTEM. This is used to specify the system ID to use for the report.

Format

SYSTEM (system ID)

Comments

- maximum of six characters
- corresponds to characters taken from .CRSID
- default = ID taken from the first record on file IF

4.4.2.3 PERIOD. This is used to specify the number of days the report will span. The PERIOD phrase has a secondary keyword followed by one or more optional parameters.

Format

```
PERIOD    DAILY   ([MMDDYY])
PERIOD    WEEKLY  ([MMDDYY])
PERIOD    MONTHLY ([MMDDYY])
PERIOD    DAYS    (#NN#) ([MMDDYY])
PERIOD    MONTHS  (#MM#) ([MMDDYY])
PERIOD    DATES   (#MMDDYY#) ([MMDDYY])
```

Comments

- [MMDDYY] is the optional start date, except for "DATES" where it is the optional stop date

- NN, MM is the number of days, months respectively

- default is all data on the file

- daily, weekly, monthly imply 1, 7, 30 days respectively.

4.4.2.4 TIME. This is used to limit the time of day covered by the report. Times are given in system time, which may not be the same as local time.

Format

```
TIME   (HHMM) (HHMM)
```

Comments

- defaults are 0001 2359

4.4.2.5 INCLUDE. The INCLUDE phrase is used to specify which plots are to be printed and to specify the disposition of periodic/nonstandard data. The INCLUDE phrase has secondary keywords followed by parameters.

Format

```
INCLUDE   GRAPH   ([1,2,...8])
INCLUDE   DATA    (#P#)
INCLUDE   DATA    (#C#)
```

Comments

- correspondence between the graph numbers and graph names is

NUMBER	GRAPH NAME
1	PROCESSOR UTILIZATION
2	MEMORY UTILIZATION
3	I/O CHANNEL TIME
4	DISK SPACE SUMMARY
5	PERIPHERAL SUMMARY
6	SCHEDULED JOBS
7	a) EXECUTION/SWAPPED SUMMARY
	b) WAITING MEMORY/PERIPHERALS
8	REMOTE STATUS
Default is	1, 2, ... 8

- P implies that data marked as periodic will be included in the report. Refer to SUMMARY statement. The default is to exclude this data.

- C implies that data marked as nonstandard configuration will be included in the report. Refer to the CONFIG statement. The default is to exclude this data.

4.4.2.6 EXCLUDE. The EXCLUDE phrase specifies plots and data to be omitted from the default list. The EXCLUDE phrase has secondary keywords followed by parameters.

Format

```
EXCLUDE GRAPH   ([1,2,...8])
EXCLUDE DATA    (#P#)
EXCLUDE DATA    (#C#)
```

Comments

- any plot numbers supplied will be excluded from the output report

- P implies that data marked as periodic will be excluded from the report

- C implies that data marked as nonstandard configuration will be excluded from the report.

4.4.2.7 Sample Parameter File. The following are three sample report statements:

```
REPORT

       NAME WEEKLY
       SYSTEM NMCC2
       PERIOD WEEKLY   082382
       TIME   1230 2030
       INCLUDE GRAPH 1,2,8

REPORT

       NAME MONTH
       SYSTEM NMCC2
       PERIOD MONTHLY  080182
       INCLUDE DATA P

REPORT

       NAME HISTRY
```

Comments

A - The weekly report will cover 7 days' data beginning on 082382 between the hours of 1230-2030. Three plots will be printed: processor utilization, memory utilization, and remote status.

B - The monitor report will cover 30 days' data beginning on 080182 between 0001-2359. Periodic data will be included. All plots will be printed.

C - The history report will summarize data for the first system ID found on the input file. Periodic data will not be included.

4.4.3 CONFIG Statement Keywords and Parameters. The keywords for the CONFIG statement are PROCESSORS, MEMORY, TAPES, DISKS, PRINTERS and CARD READERS. The CONFIG statement is processed by PSUMR to determine when data is collected under standard/nonstandard conditions. If a CONFIG statement is present, the variables will be checked against the system configuration. If the system configuration is not the same as the CONFIG statement, the intermediate record will be flagged.

4.4.3.1 PROCESSOR. PROCESSOR specifies the standard number of processors.

Format

```
PROCESSORS  (#Y#)
PROC  (#Y#)
```

Comments

- y = number processors

4.4.3.2 MEMORY. Specifies the standard memory size in K.

Format

MEMORY (#YYYY#)

Comments

- YYYY must be divisable by 64.

4.4.3.3 TAPES. TAPES specifies the standard number of physical channels which have a tape controller configured and the number of tape drives.

Format

TAPE (#X/YY#)
TAPES (#X/YY#)

Comments

- X = the number of physical IOM channels

- YY = the number of tape drives

- only one TAPE phrase per CONFIG statement is allowed. Thus, X, YY should be the totals for the system.

4.4.3.4 DISKS. Specifies the device type, number of physical channels, and number of devices for disks. There should be one DISK phase for each device class configured.

Format

DISK (#DDD/X/YY#)
DISKS (#DDD/X/YY#)

Comments

- DDD = the device type. Permissable device types are: 167, 170, 180, 181, 190, 191, 310, 400, 450, 500, 501

- X = the number of physical IOM channels for this device

- YY = the number of disk drives of this type.

4.4.3.5 PRINTERS. Specifies the standard number of printers configured.

Format

```
PRINTERS    (#X#)
PRINT       (#X#)
```

Comments

- X = the number of printers

4.4.3.6 Sample Parameter File. The following is a sample configuration statement:

```
CONFIG

      PROC 2, MEMORY 768

      TAPES  3/18
      DISKS  191/1/19
      DISKS  450/2/27
      PRINT  4
```

4.4.4 SUMMARY Statement Keywords and Parameters. The keywords for the SUMMARY statement are NOSUM, SUM. This statement is used to indicate what data is to be included/excluded from standard reports.

4.4.4.1 NOSUM. NOSUM specifies a sequence of intervals for which data will be included or omitted (marked "periodic").

Format

```
NOSUM  (#MMDDYY/I/O#)
```

Comments

- I = the number of days to include
- O = the number of days to exclude

Example

```
NOSUM  082382/5/2
```

This statement says "beginning with 23 August 1982, include 5 days' data, then omit 2 days." The initial include/omit intervals would be

INCLUDE	OMIT
8/23/82-8/27/82	8/28/82-8/29/82
8/30/82-9/3/82	9/4/82-9/5/82
9/6/82-9/10/82	9/11/82-9/12/82

4.4.4.2 SUM. This specifies the time of day for which data will be included or omitted (marked "periodic").

Format

```
SUM  (#HHMM HHMM#)
```

Comments

- This gives the start/stop interval for including data. Times which are not within the interval will be omitted from standard summaries.

4.4.5 SNUMB Statement Keywords and Parameters. The keywords for the SNUMB statement are SYSTEM and SPECIAL. The purpose of this statement is to increment the default list of system SNUMBs and to allow the user to select SNUMBs to be plotted separately using the letters A, B on the graphs.

4.4.5.1 SYSTEM. This specifies SNUMBs which are to be added to the default list of system jobs. The default list is shown in figure 4-2.

Format

```
SYSTEM  (#AAAAA,BBBBB,...#)
```

Comments

- the data for the given SNUMBs will be plotted with system SNUMBs.

4.4.5.2 SPECIAL. This specifies SNUMBs which will be plotted separately using the letters A or B.

Format

```
SPECIAL  A  (#AAAAA,BBBBB,...#)
SPECIAL  B  (#XXXXX,YYYYY,...#)
```

Comments

- the data for AAAAA,BBBBB which appear on the plots using the letter A, while XXXXX,YYYYY will be plotted using the letter B.

4.4.5.3 Sample SNUMB Statement. The following is a sample SNUMB statement:

```
SNUMBS  SYSTEM  NACE,DMTEX,DMSTA

SPECIAL  A  [illegible]  [illegible]  FTS  TLCF

[illegible]  [illegible]  [illegible]  [illegible]NP-C
```

$CALC
$PALC
$SYOT
$RTIN
$TDIS
$FSYS
COLCT
HEALS
VIDEO
$IMCV
$PACT
$TRAX
$MOLT
$POLT
$COLT
$TOLT
$SOLT
$SLTA
GEIN

Figure 4-2. Default System SNUMBs

4.5 Description of Reports

Each report produced by RPTSUM or DEMON consists of a banner page, report overview page, and one or more graphs.

4.5.1 Report Overview Page. (See figure 4-3). This page lists the report identification, system ID, dates, times and requested graphs on the left. On the right is a summmary including the number of processors, tape channels, tape drives, disk channels, disk drives, printer channels and printers. Additionally, the number of disk drives, by type, is listed.

4.5.2 Graphs. In general, the graphs are "bar" graphs which describe the utilization of one or more system components over time. The time shown at the bottom of each graph is the system time which may be different than local time. At the right hand side of the graph page is a listing/description of the codes used within the graph. Below the usage codes is a summary showing the maximum and average values for each code graphing. Below this summary, the dates and times are repeated. Finally, the horizontal (time) and increments are shown.

4.5.2.1 Usage Codes. The chart below lists the usage codes and an explanation of each:

Code	Explanation
System = S	All jobs classified as system jobs (figure 4-2) plus any jobs specified as system in a SNUMB statement
TSS = *	Timesharing utilization
"T" = T	Jobs submitted via CARDIN (i.e., SNUMB = XXXXT)
Other = O	All jobs which are not included in another catagory
User A = A	Those jobs specified as Special A in a SNUMB statement
User B = B	Those jobs specified as Special B in a SNUMB statement
GCOS HC = H	The amount of memory currently used for GCOS hard core
DISC = D	The percent of available channel time used for disk I/O
TAPE = T	The percent of available channel time used for tape I/O

```
DATE  01-11-83

DAILY 15  01-11-83   01.040          DAILY RESOURCE MONITOR                    PAGE   14

    REPORT IDENTIFICATION = DAILY                           CONFIGURATION SUMMARY      DISK  SUMMARY

    SYSTEM IDENTIFICATION = NMCC2                           PROCESSORS =   2           DEVICE    NUMBER

    INCLUSIVE DATES       =  01-10-83    01-10-83           MEMORY     =   0768K        191        27
                                                                                        450        18
    INCLUSIVE TIMES       =  00:01       23:59              TAPE CHANNELS  =  03

    GRAPH LIST                                              TAPE DRIVES    =  18

                                                            DISK CHANNELS  =  03
     NUMBER    DESCRIPTION
       1        PROCESSOR UTILIZATION                       DISK DRIVES    =  45

       2        MEMORY UTILIZATION                          PRINT CHANNELS =  06

       3        I/O CHANNEL TIME                            PRINTERS       =  06

       4        DISK SPACE SUMMARY

       5        PERIPHERAL SUMMARY

       6        SCHEDULED JOBS

       7        EXECUTION/SWAPPED SUMMARY

       8        WAITING MEMORY/PERIPHERALS

       9        REMOTE STATUS
```

Figure 4-3. Daily Resource Monitor Overview

USERID OPNSUTIL

PRM Avail = A	The number of links in thousands of available permanent file space
PRM Assg = P	The number of links in thousands of catalogued permanent file space
RMB Avail = B	The number of links in thousands of available space on removable disk packs
RMB Assg = N	The number of links in thousands of catalogued space on removable disk packs
Defective = D	The number of links in thousands of disk space marked defective
Assigned = A	The number of devices assigned
Released = R	The number of devices released
Dedicated = D	The number of devices dedicated
Available = F	The number of devices available for assignment
Express = E	The number of jobs found in the express queue
HOLD = H	The number of jobs found in the hold queue
CLOSED = C	The number of scheduler classes marked closed
TSS = T	The number of users currently logged onto Timesharing
REMOTE = R	The number of lines logged onto the system other than those logged onto Timesharing

4.5.2.2 Report Graphs. The graphs available for reports are:

Graph Number	Title	Figure Number
1	Processor Utilization	4-4
2	Memory Utilization	4-5
3	I/O Channel Time	4-6
4	Disk Space Summary	4-7
5	Peripheral Summary	4-8
6	Scheduled Jobs	4-9
7	(a) Execution/Swapped Summary	4-10
	(b) Waiting Memory/Peripherals	4-11
8	Remote Status	4-12

DATE 01-11-83 UNCLASSIFIED

DAILY 15 01-11-83 01.040 DAILY RESOURCE MONITOR PAGE 16

1 PROCESSOR UTILIZATION

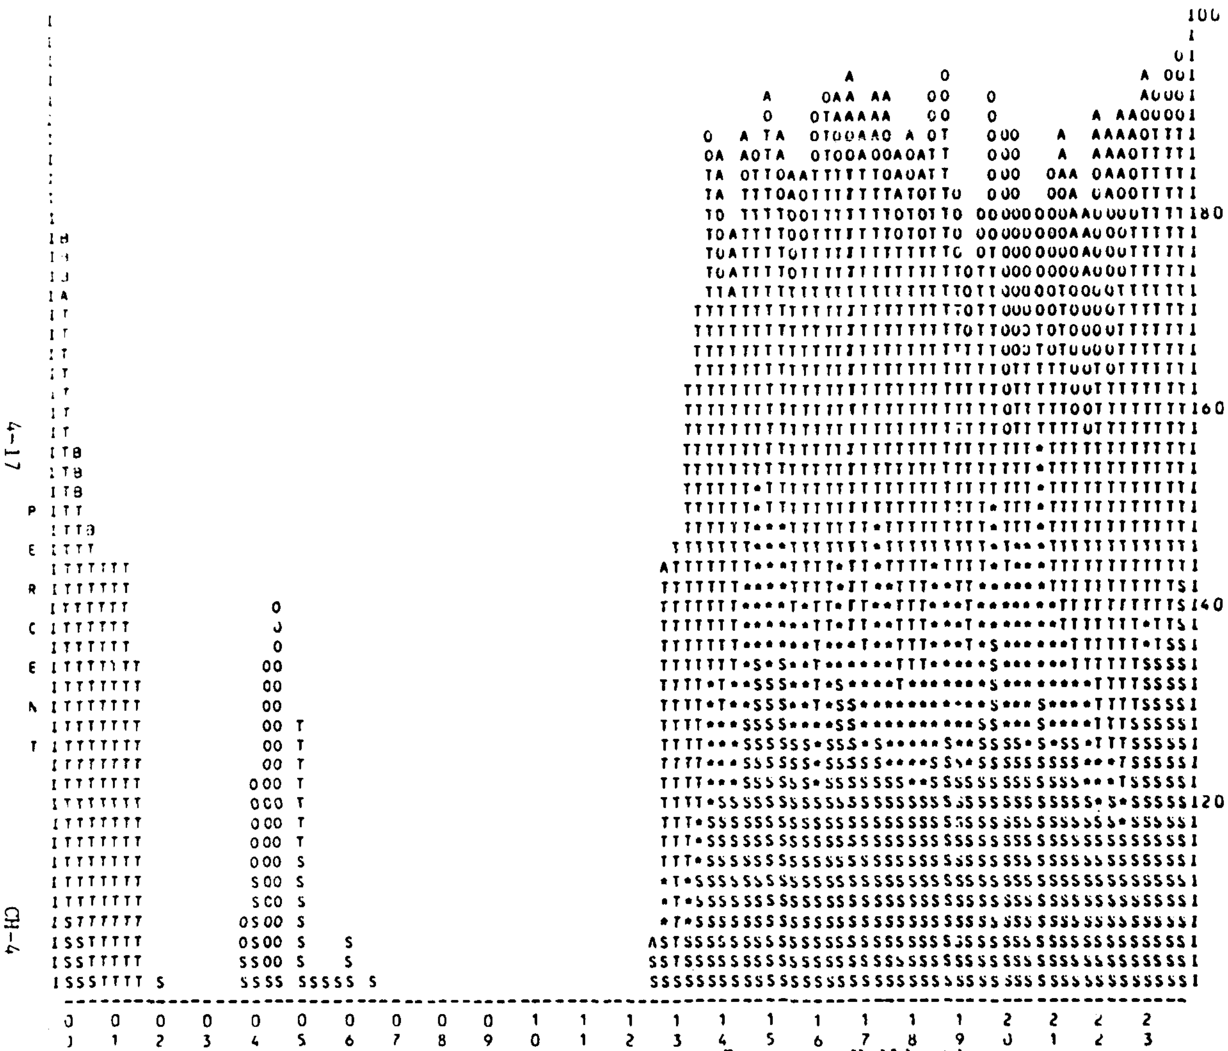

USAGE CODE

SYSTEM = S
TSS = *
"T" = T
OTHER = O
USER A = A
USER B = B

** SUMMARY **

TYPE	MAX	AVG
S	42	17
*	26	12
T	63	40
O	38	10
A	9	4
B	6	5

	FROM	TO
DATE	011063	011063
TIME	00:01	23:59

TIME INCREMENT
= 15 MIN
PER CENT INCR
= 2%

USERID OPNSUTIL

Figure 4-4. Processor Utilization

DATE 01-11-83 UNCLASSIFIED

DAILY 15 01-11-83 01.040 DAILY RESOURCE MONITOR PAGE 15

2 MEMORY UTILIZATION

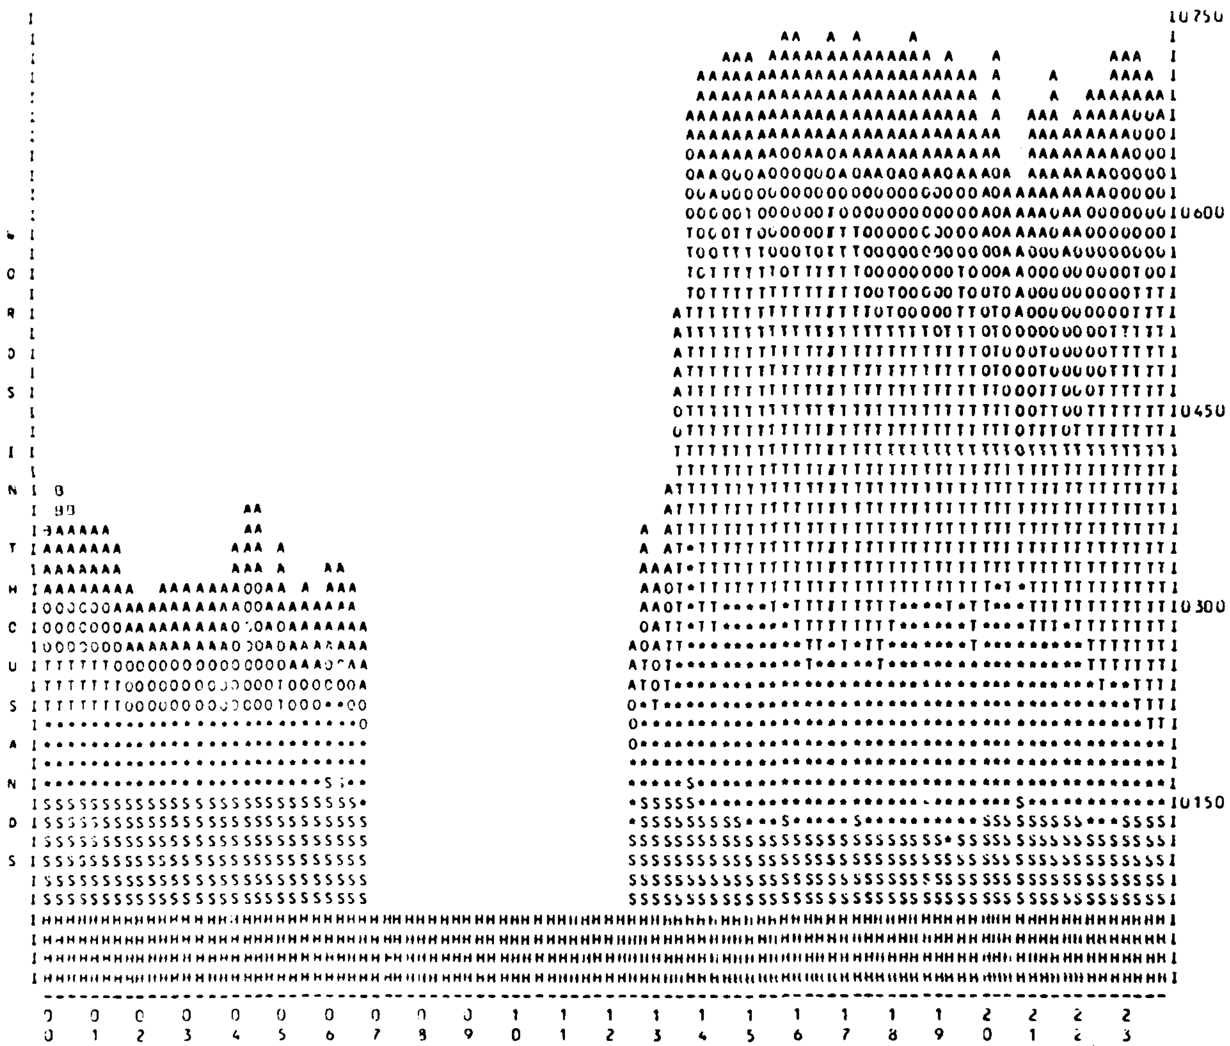

USAGE CODE

GCOS HC = H
SYSTEM = S
TSS = *
"T" = T
OTHER = O
USER A = A
USER B = B

** SUMMARY **

TYPE	MAX	AVG
H		54
S	104	77
*	186	110
T	350	196
O	150	71
A	111	77
B	30	7

	FROM	TO
DATE	011083	011083
TIME	00:01	23:59

TIME INCREMENT
= 15 MIN
WORDS INCR
= 15K

USERID OPNSUTIL

Figure 4-5. Memory Utilization

DATE 01-11-83 UNCLASSIFIED

DAILY 15 01-11-83 01.040 DAILY RESOURCE MONITOR PAGE 17

3 I/O CHANNEL TIME

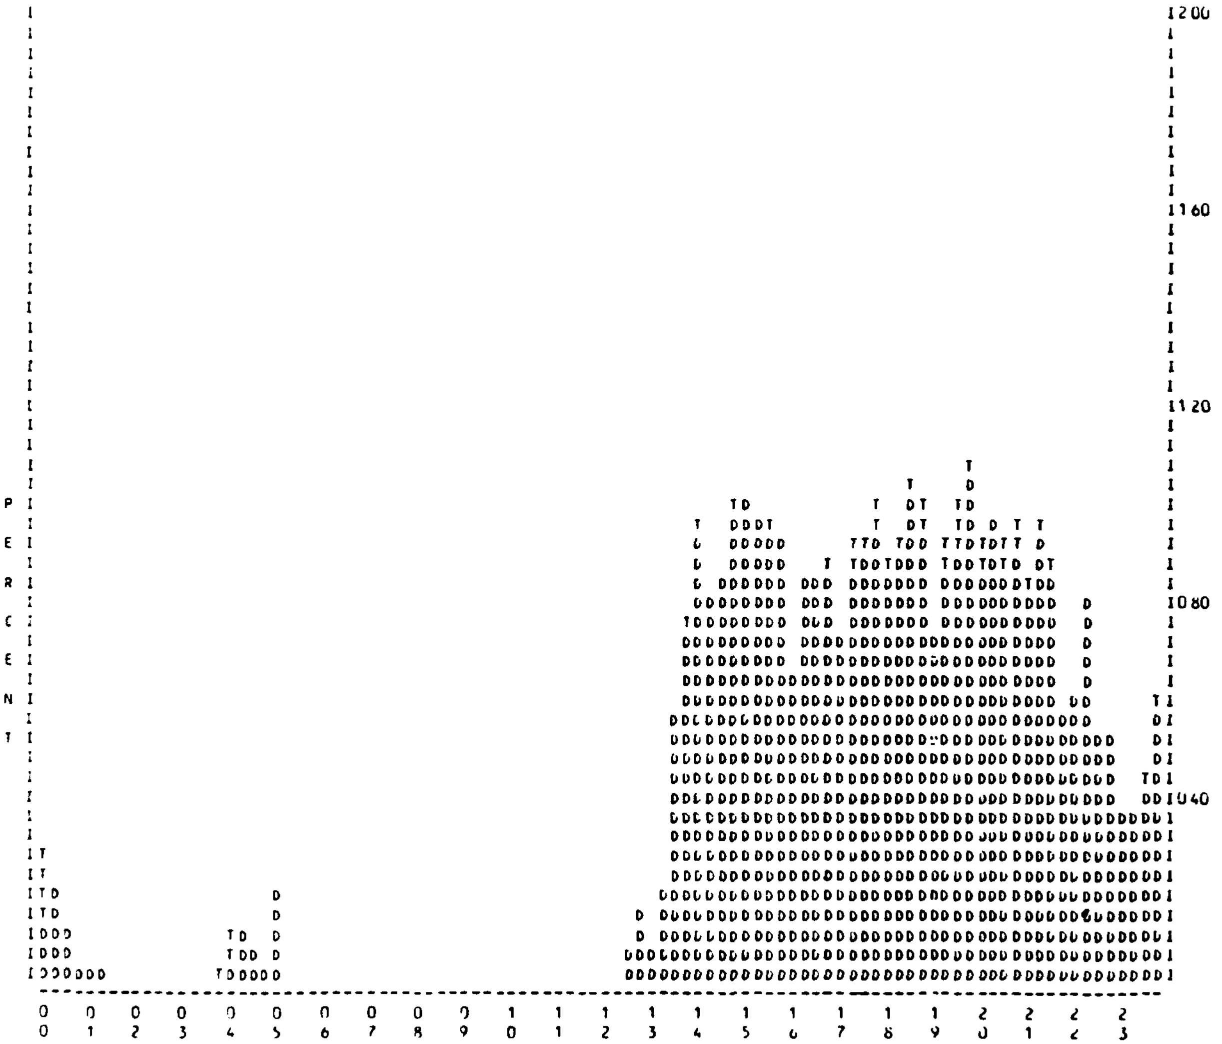

USAGE CODE

DISC = D
TAPE = T
PRINT = P

** SUMMARY **

TYPE	MAX	AVG
D	105	53
T	18	4
P		

	FROM	TO
DATE	011063	011063
TIME	00:01	23:59

TIME INCREMENT
= 15
PER CENT INCR
= 4%

USERID OPNSUTIL

Figure 4-6. I/O Channel Time

DATE 01-11-83 UNCLASSIFIED

DAILY 15 01-11-83 01.040 DAILY RESOURCE MONITOR PAGE 18

4 DISK SPACE SUMMARY

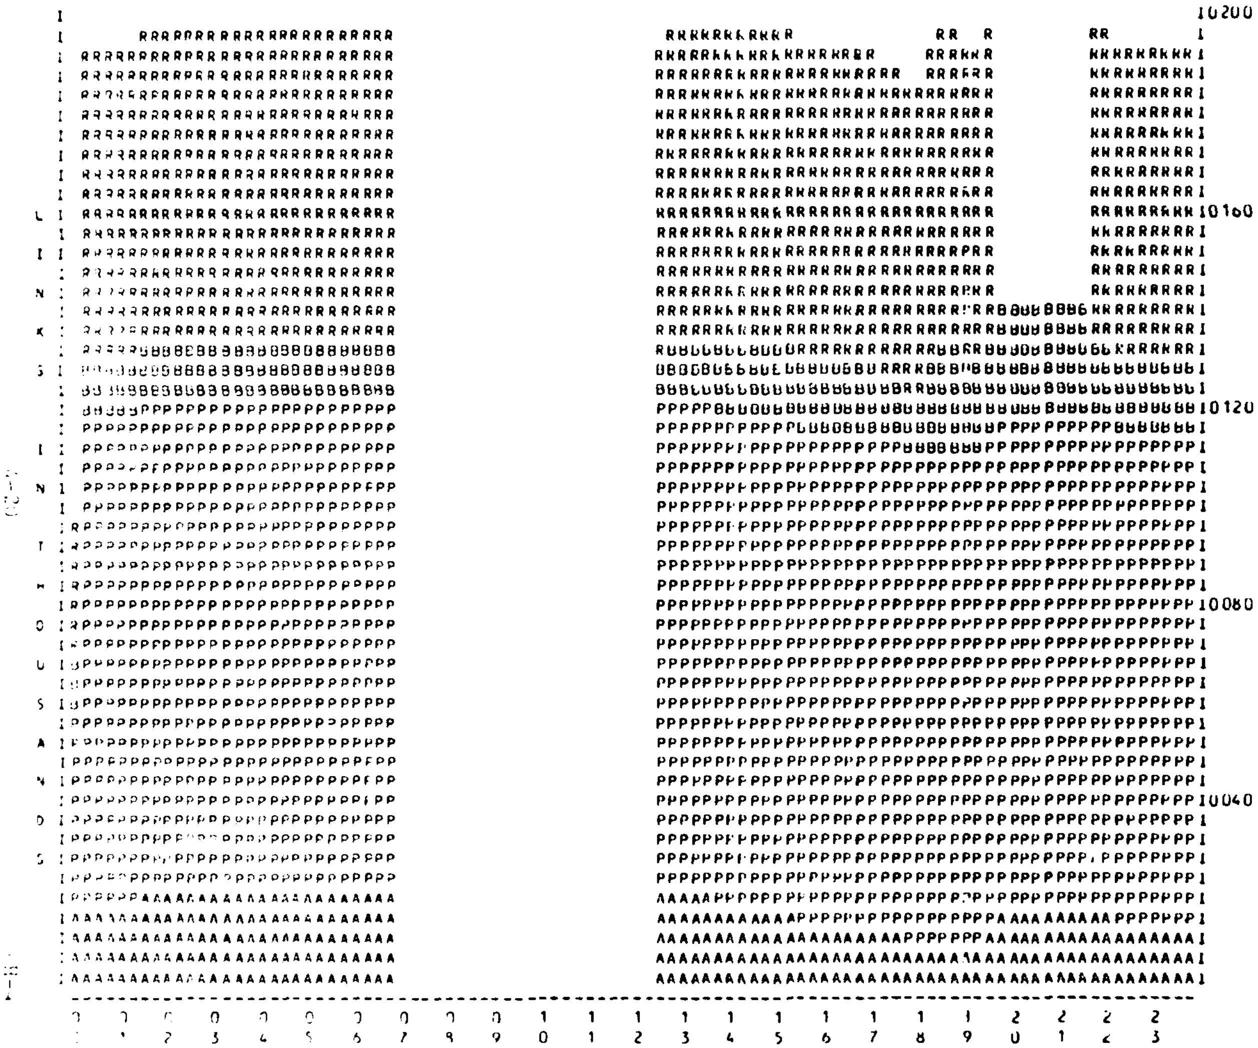

USAGE CODE

PRM AVAIL = A
PRM ASSG = P
RMB AVAIL = B
RMB ASSG = R
DEFECTIVE = D

** SUMMARY **

TYPE	MAX	AVG
A	22	17
P	100	99
B	23	16
R	65	64
D		.006

	FROM	TO
DATE	011083	011083
TIME	00:01	23:59

TIME INCREMENT
= 15 MIN
LINKS INCREMENT
= 0004K LINKS

Figure 4-7. Disk Space Summary

USERID OPNSUTIL

```
DATE  01-11-83                         UNCLASSIFIED
DAILY 15  01-11-83   01.040        DAILY RESOURCE MONITOR                          PAGE     19
                                   5    PERIPHERAL SUMMARY                                  USAGE  CODE

                                                                                            ASSIGNED   = A
                                                                                            RELEASED   = R
                                                                                            DEDICATED  = D
                                                                                            AVAILABLE  = F
```

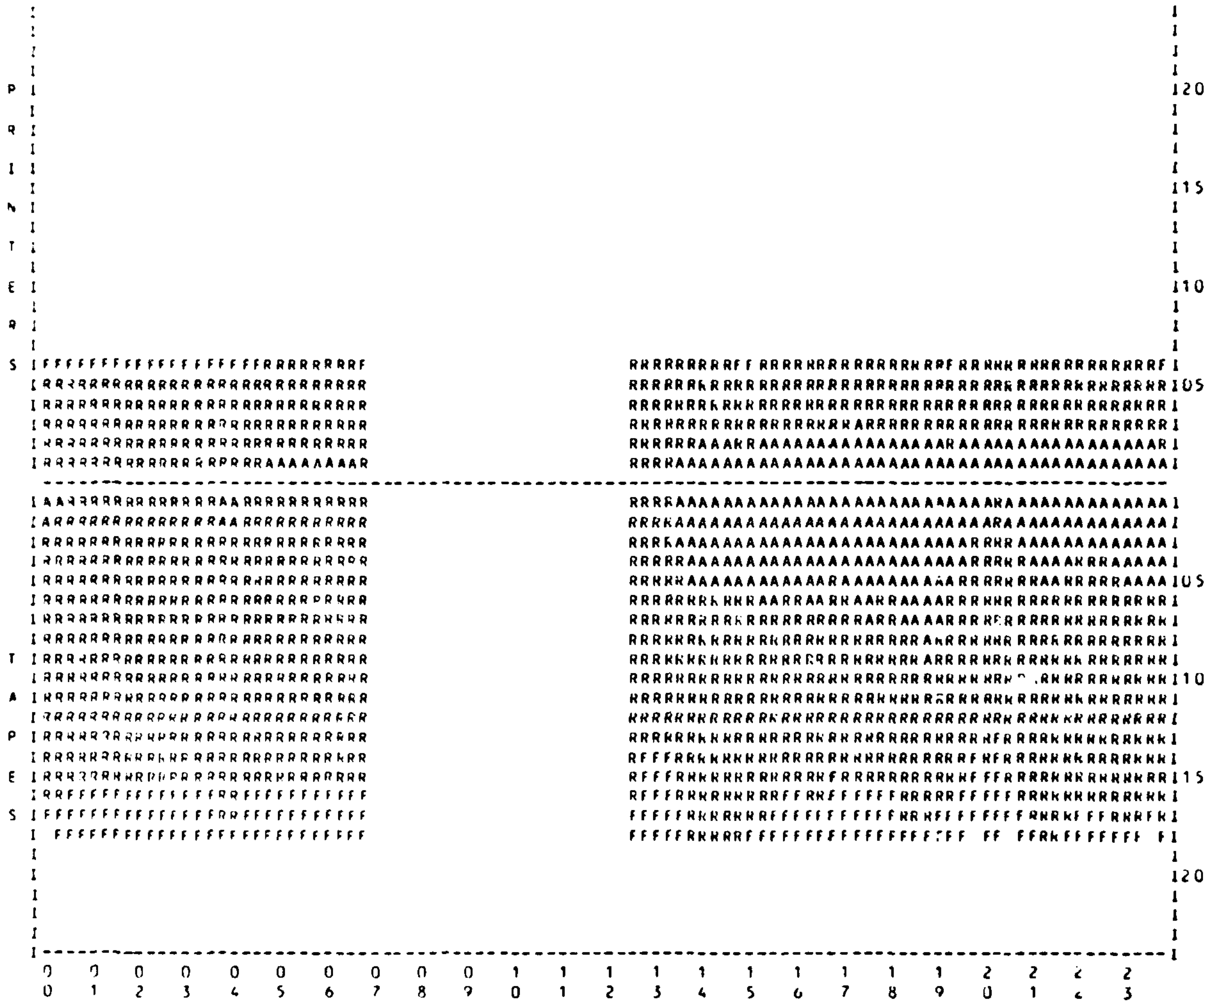

```
** SUMMARY **

     PRINTERS
TYPE  MAX  AVG

 A     3    1
 R     0    4
 D
 F     1    1

      TAPES
TYPE  MAX  AVG

 A     9    3
 R    16   13
 D
 F     6    3

        FROM      TO
DATE   011083   011063

TIME   00:01    23:59

TIME INCREMENT
   = 15  MIN
Y AXIS INCREMENT
PRT = 01
TAP = 01

USERID  OPNSUTIL
```

Figure 4-8. Peripheral Summary

DATE 01-11-83 UNCLASSIFIED

DAILY 15 01-11-83 01.040 . DAILY RESOURCE MONITOR PAGE 20

6 SCHEDULED JOBS

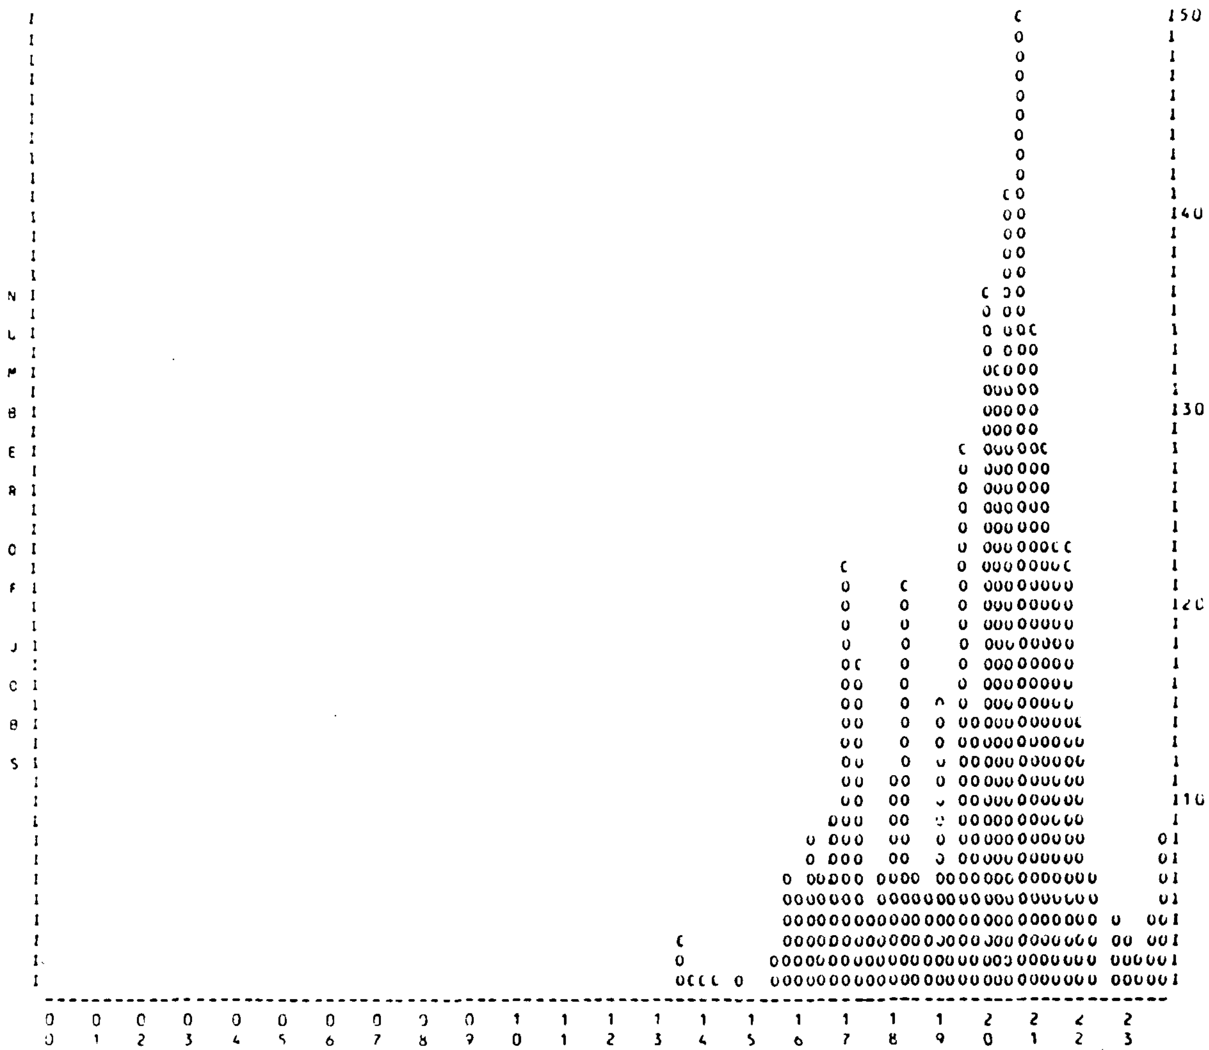

USAGE CODE

EXPRESS = E
HOLD = H
OTHER = O
CLOSED = C

** SUMMARY **

TYPE	MAX	AVG
E		
H		
O	49	13
C	2	

	FROM	TO
DATE	011083	011083
TIME	00:01	23:59

TIME INCREMENT
= 15 MIN
JOB INCR
= 01

USERID OPNSUTIL

Figure 4-9. Scheduled Jobs

DATE 01-11-83 UNCLASSIFIED

DAILY 15 01-11-83 01.040 DAILY RESOURCE MONITOR PAGE 21

7 EXECUTION/SWAPPED SUMMARY

USAGE CODE

SYSTEM	=	S
TSS	=	*
"T"	=	T
OTHER	=	O
USER A	=	A
USER B	=	B

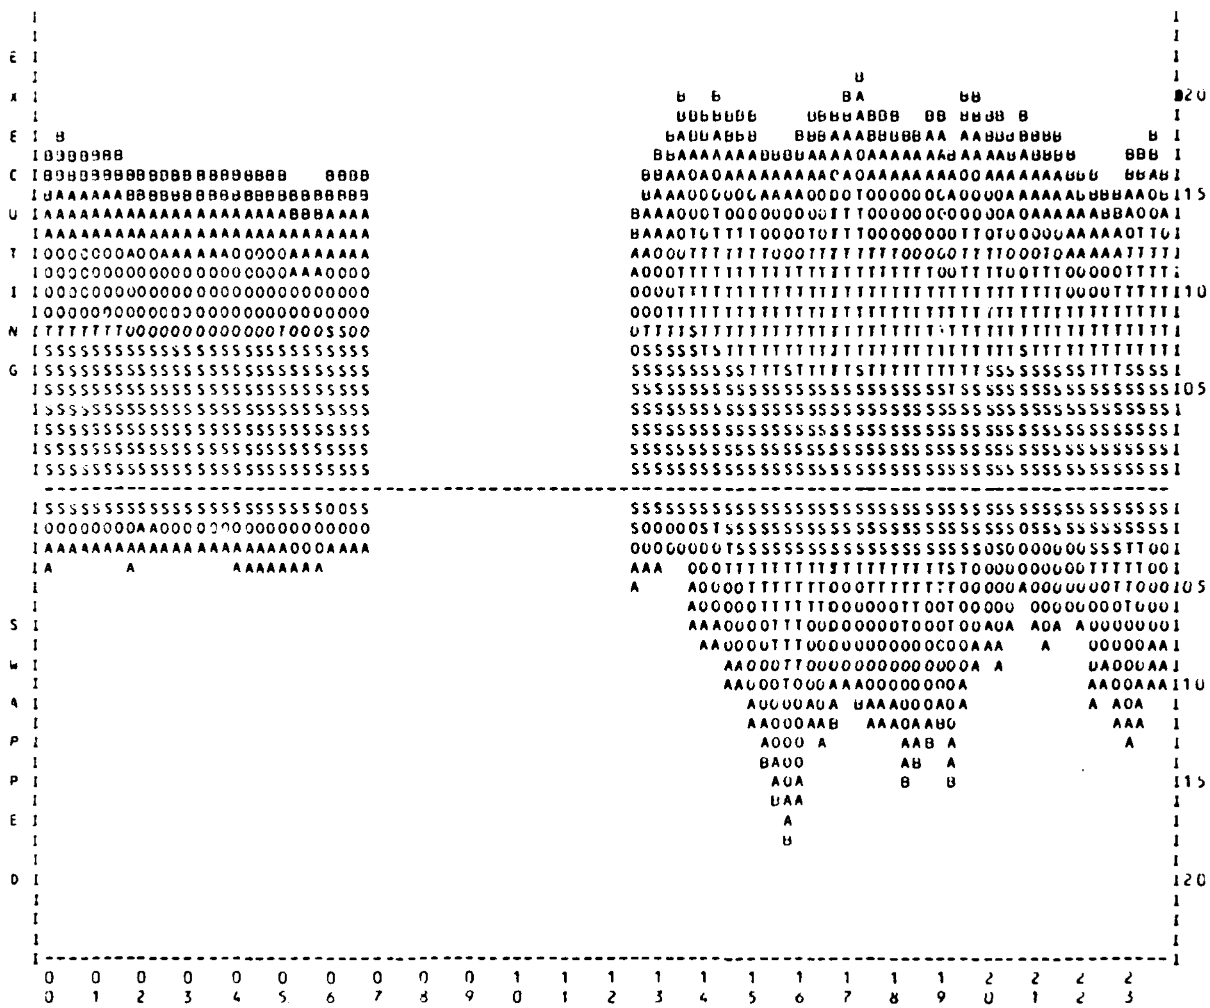

** SUMMARY **
** IN EXECUTION **

TYPE	MAX	AVG
S	8	6
*		
T	9	5
O	5	3
A	4	2
B	3	2

** SWAPPED **

TYPE	MAX	AVG
S	4	2
*		
T	7	1
O	6	3
A	3	2
B	1	

	FROM	TO
DATE	011083	011083
TIME	00:01	23:59

TIME INCREMENT
= 15 MIN
Y AXIS INCR
EXEC = 01
SWAP = 01

USERID OPNSUTIL

Figure 4-10. Execution/Swapped Summary

DATE 01-11-83 UNCLASSIFIED

DAILY 15 01-11-83 01.040 DAILY RESOURCE MONITOR PAGE 22

8 WAITING MEMORY/PERIPHERALS

USAGE CODE

SYSTEM = S
TSS = *
"T" = T
OTHER = O
USER A = A
USER B = B

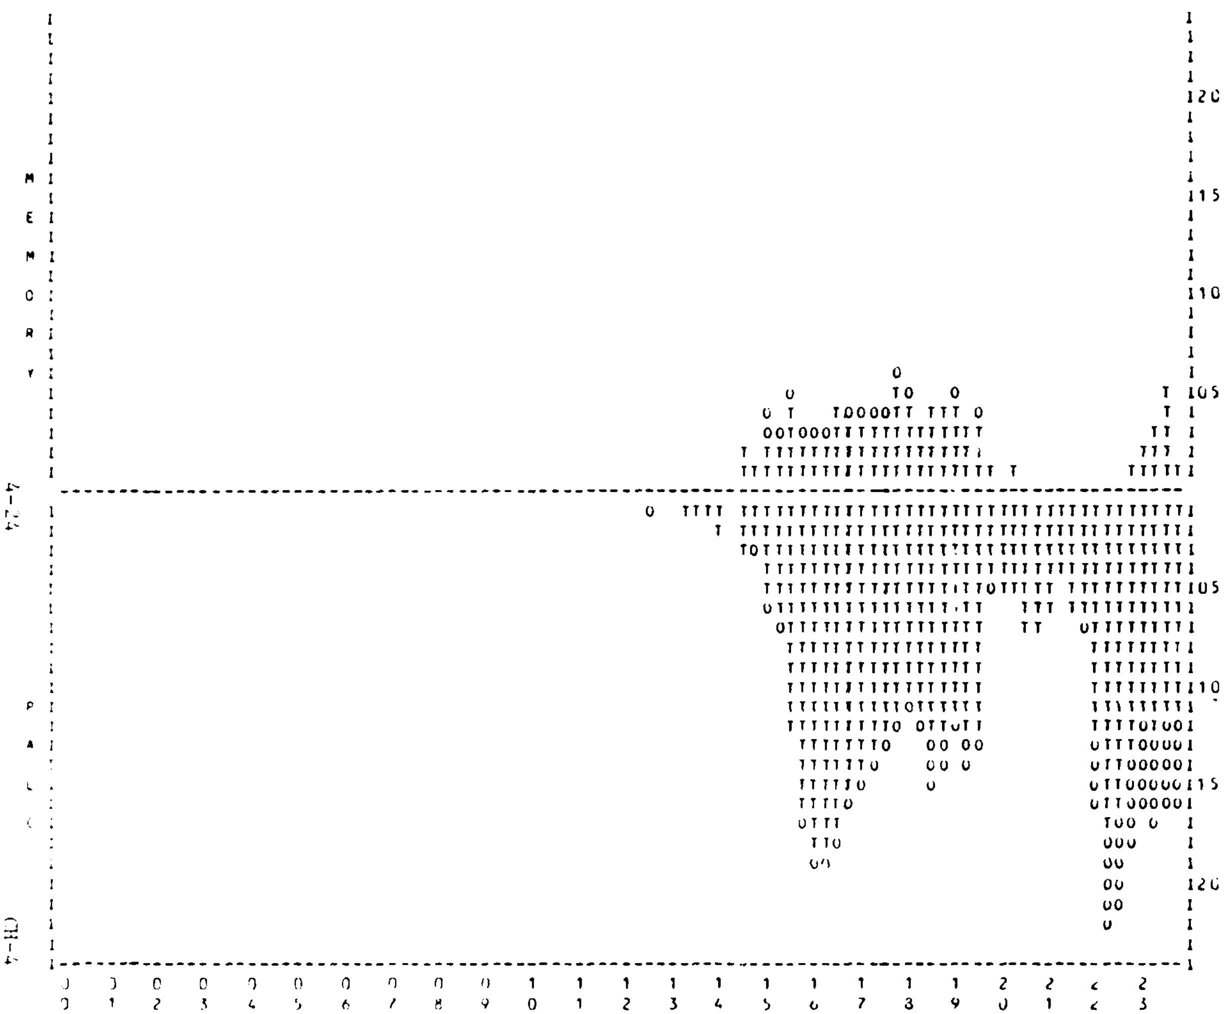

** SUMMARY **

** WAITING MEMORY **

TYPE	MAX	AVG
S		
*		
T	5	2
O	2	
A		
B		

** WAITING PALC **

TYPE	MAX	AVG
S		
*		
T	18	10
O	5	2
A		
B		

	FROM	TO
DATE	011083	011083
TIME	00:01	23:59

TIME INCREMENT = 15 MIN
Y AXIS INCR
EXEC = 01
MEM = 01

USERID OPNSUTIL

Figure 4-31. Waiting Memory/Peripherals

DATE 01-11-83 UNCLASSIFIED

DAILY 15 01-11-83 01.040 DAILY RESOURCE MONITOR PAGE 23

9 REMOTE STATUS

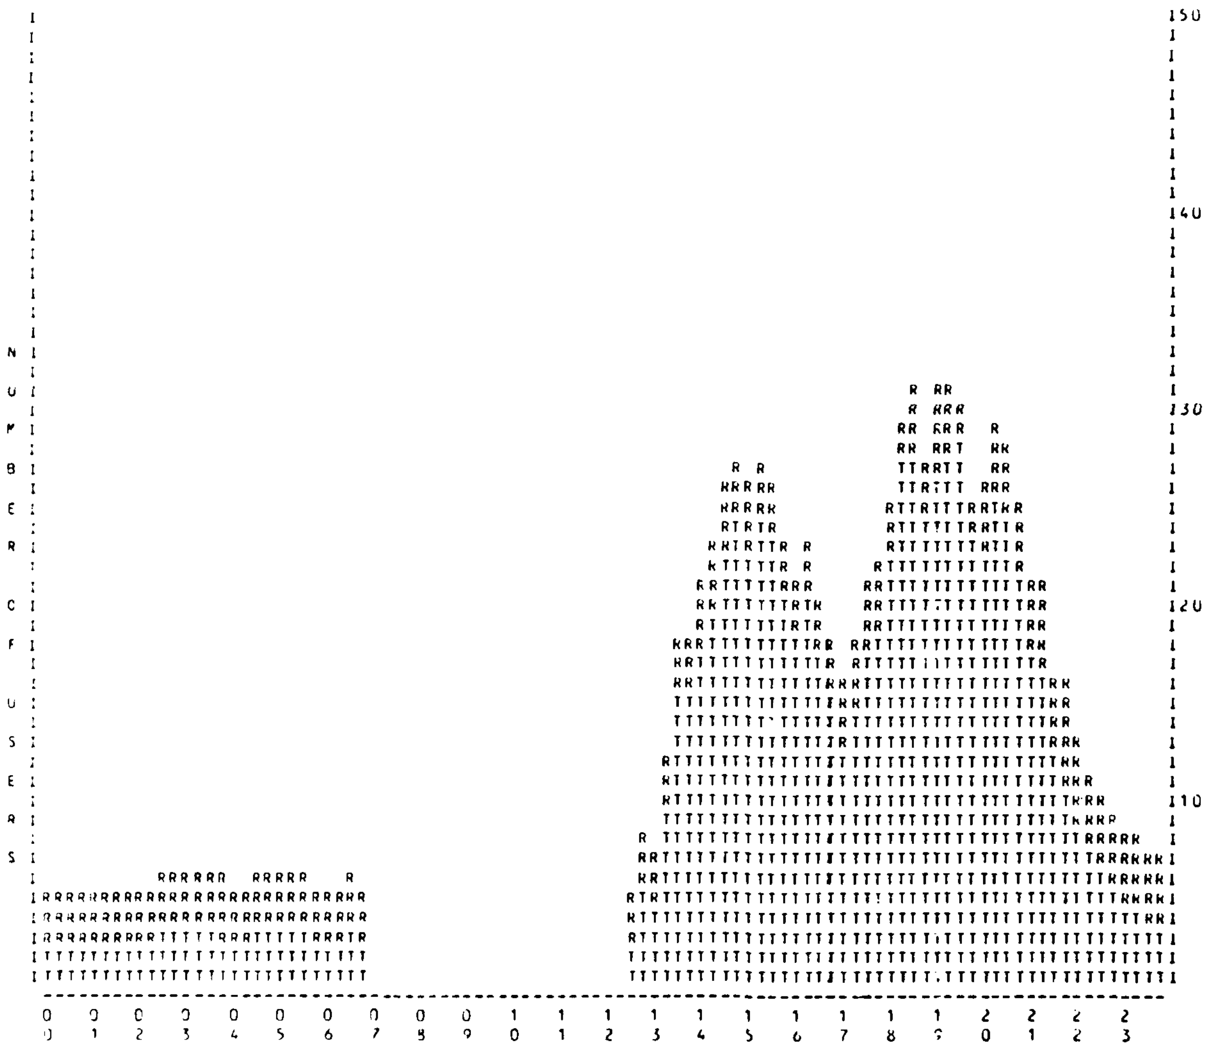

USAGE CODE

TSS = T
REMOTE = R

** SUMMARY **

TYPE	MAX	AVG
T	28	11
R	6	3

	FROM	TO
DATE	011083	011083
TIME	00:01	23:59

TIME INCREMENT
= 15 MIN
USERS INCREMENT
= 01

Figure 4-12. Remote Status

USERID OPNSUTIL

4.6 Usage Concepts

The WWMCCS resource monitor is designed to provide a means of tracking the utilization of major system components over extended time periods and to daily monitor the system usage automatically. By using the data supplied from the graphs, a PLOT of the system utilization can be developed. Such a plot would provide useful trending information. The daily monitoring capability provides a means of determining when utilization has varied from the established norm.

To use the resource monitor as designed, the following steps are required:

1. Turn on collection of SCF 608 records (refer to section 3).

2. Run the collector while the system is operational.

3. At the end of the system clock day (2400Z for WWMCCS sites), close the SCF file.

4. Run a PSUMR/DEMON job to create intermediate records and update the history file. The input is the day's SCFCLO tapes (or a daily tape made from the SCFCLO tapes). The intermediate records should be accumulated on a disk file if a weekly report or tending analysis will be done. The parameters of the history file should be set for prime time. This is accomplished by having the PSUMR/DEMON parameter file contain the following SUMMARY statement:

```
SUMMARY
        SUM     1300 2100
        NOSUM   011083/5/2
```

   Washington is ZULU +5 so prime time is 0800 + 5 = 1300. Hawaii is +10; hence, prime time begins at 1800. The NOSUM statement will exclude weekend data from the history file.

5. Daily Monitoring - Once the history file is initialzed, the graphs will not be generated unless the current day's data differs from the history file data by 10 percent. When the graphs are generated, it is an "automatic" signal that this day's utilization varied from the norm. The current status of the history file can be obtained (i.e., a set of graphs generated) by running DEMON and omitting file code for input.

6. Trending Analysis - In order to detect trends (increases or decreases) in utilization, the analyst must develop a plot of the values he wishes to "track." Until more experience is gained with the system, we recommend a weekly plot because a daily plot will probably have to much variance, while monthly might not be sufficient to detect trends. The weekly figures to plot are

obtained by generating desired graphs using RPTSUM and the saved intermediated record file. For example: suppose the site wishes to track TSS processor utilization and the number of TSS users during prime and non-prime time. This would be accomplished by running RPTSUM with the following parameters (assume +5 for time):

```
REPORT

     NAME PRIME
     SYSTEM NMCC
     TIME   1300 2100
     INCLUDG GRAPH 1,8

REPORT

     NAME NPRIME
     SYSTEM NMCC
     TIME   2101 1259
     INCLUDE GRAPH 1,8
```

These parameters will generate two processor utilization graphs and two remote status graphs: one for prime and one for non-prime time. Using graph paper, record the average (or maximum) TSS processor value and number of TSS users (see figure 4-13). If a site is currently experiencing a change in utilization, several weeks' data should show a trend (increase/decrease).

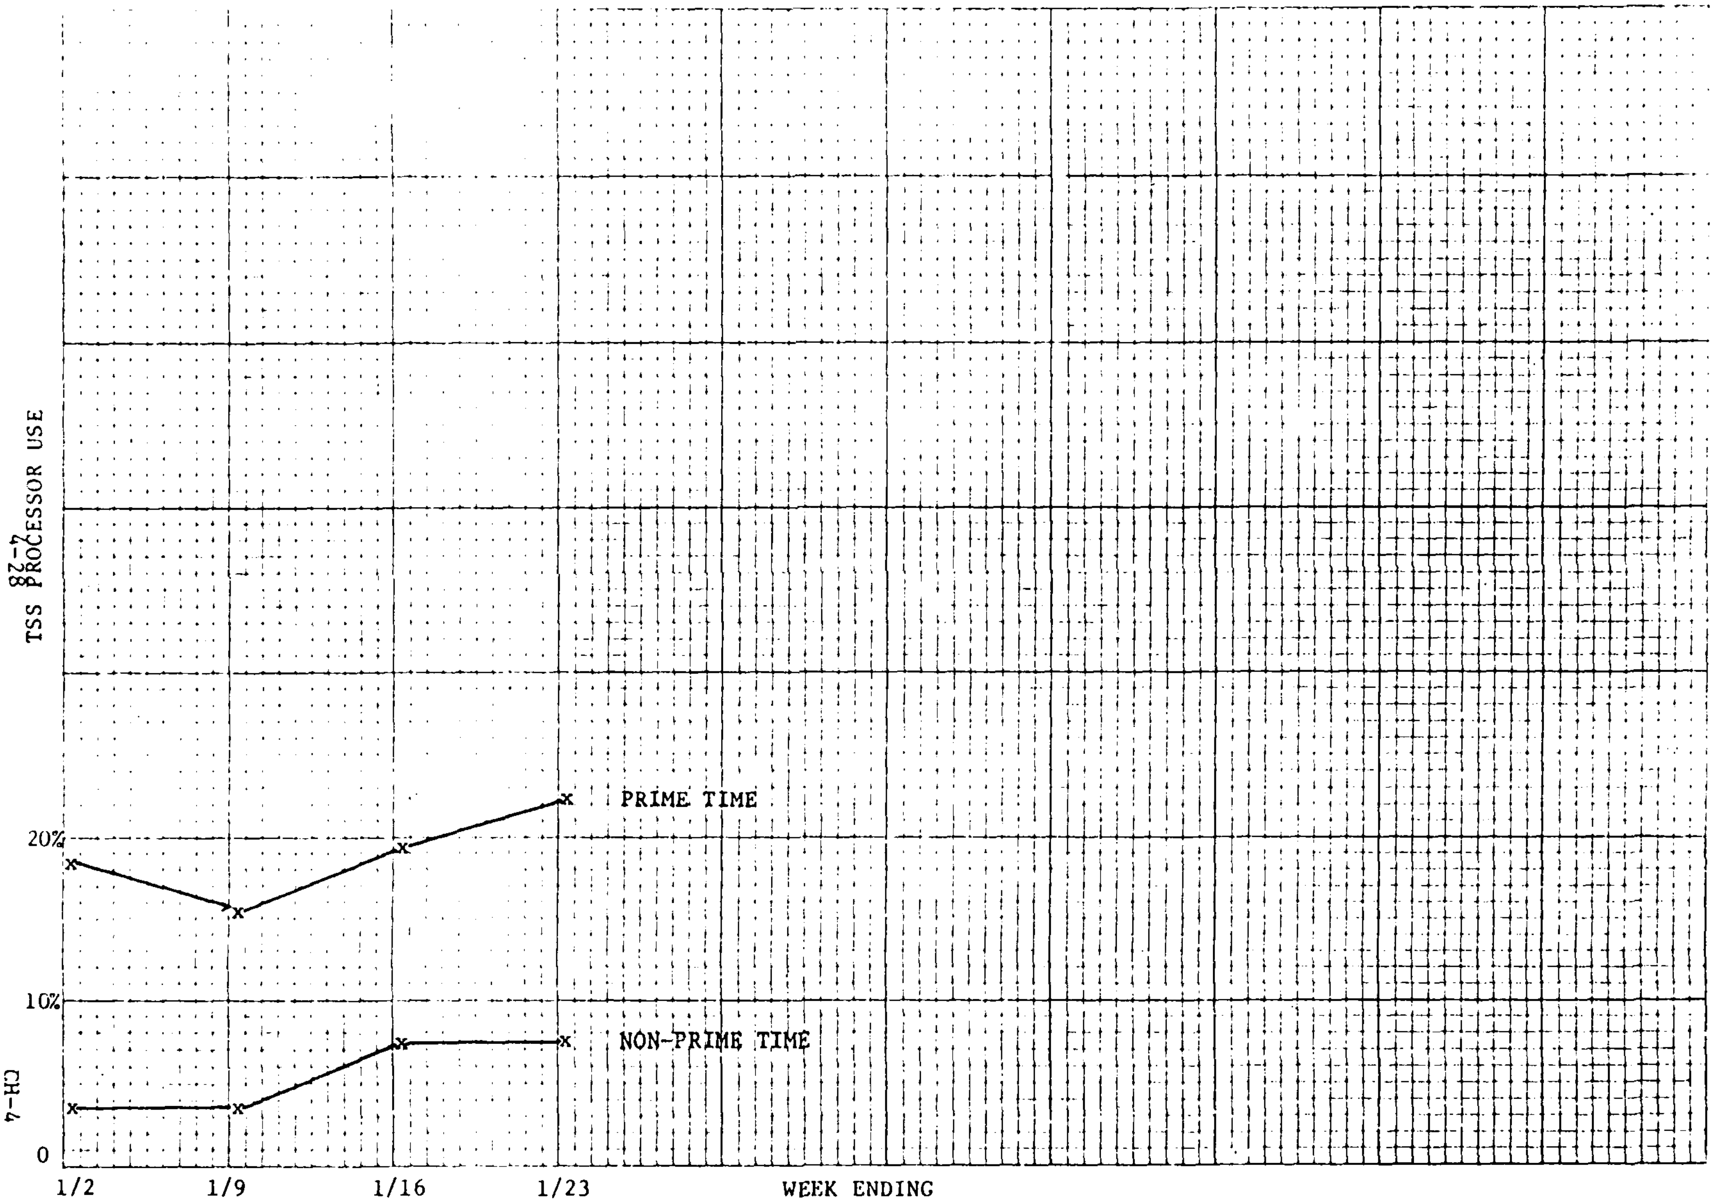

Figure 4-13. TSS Usage (Part 1 of 2)

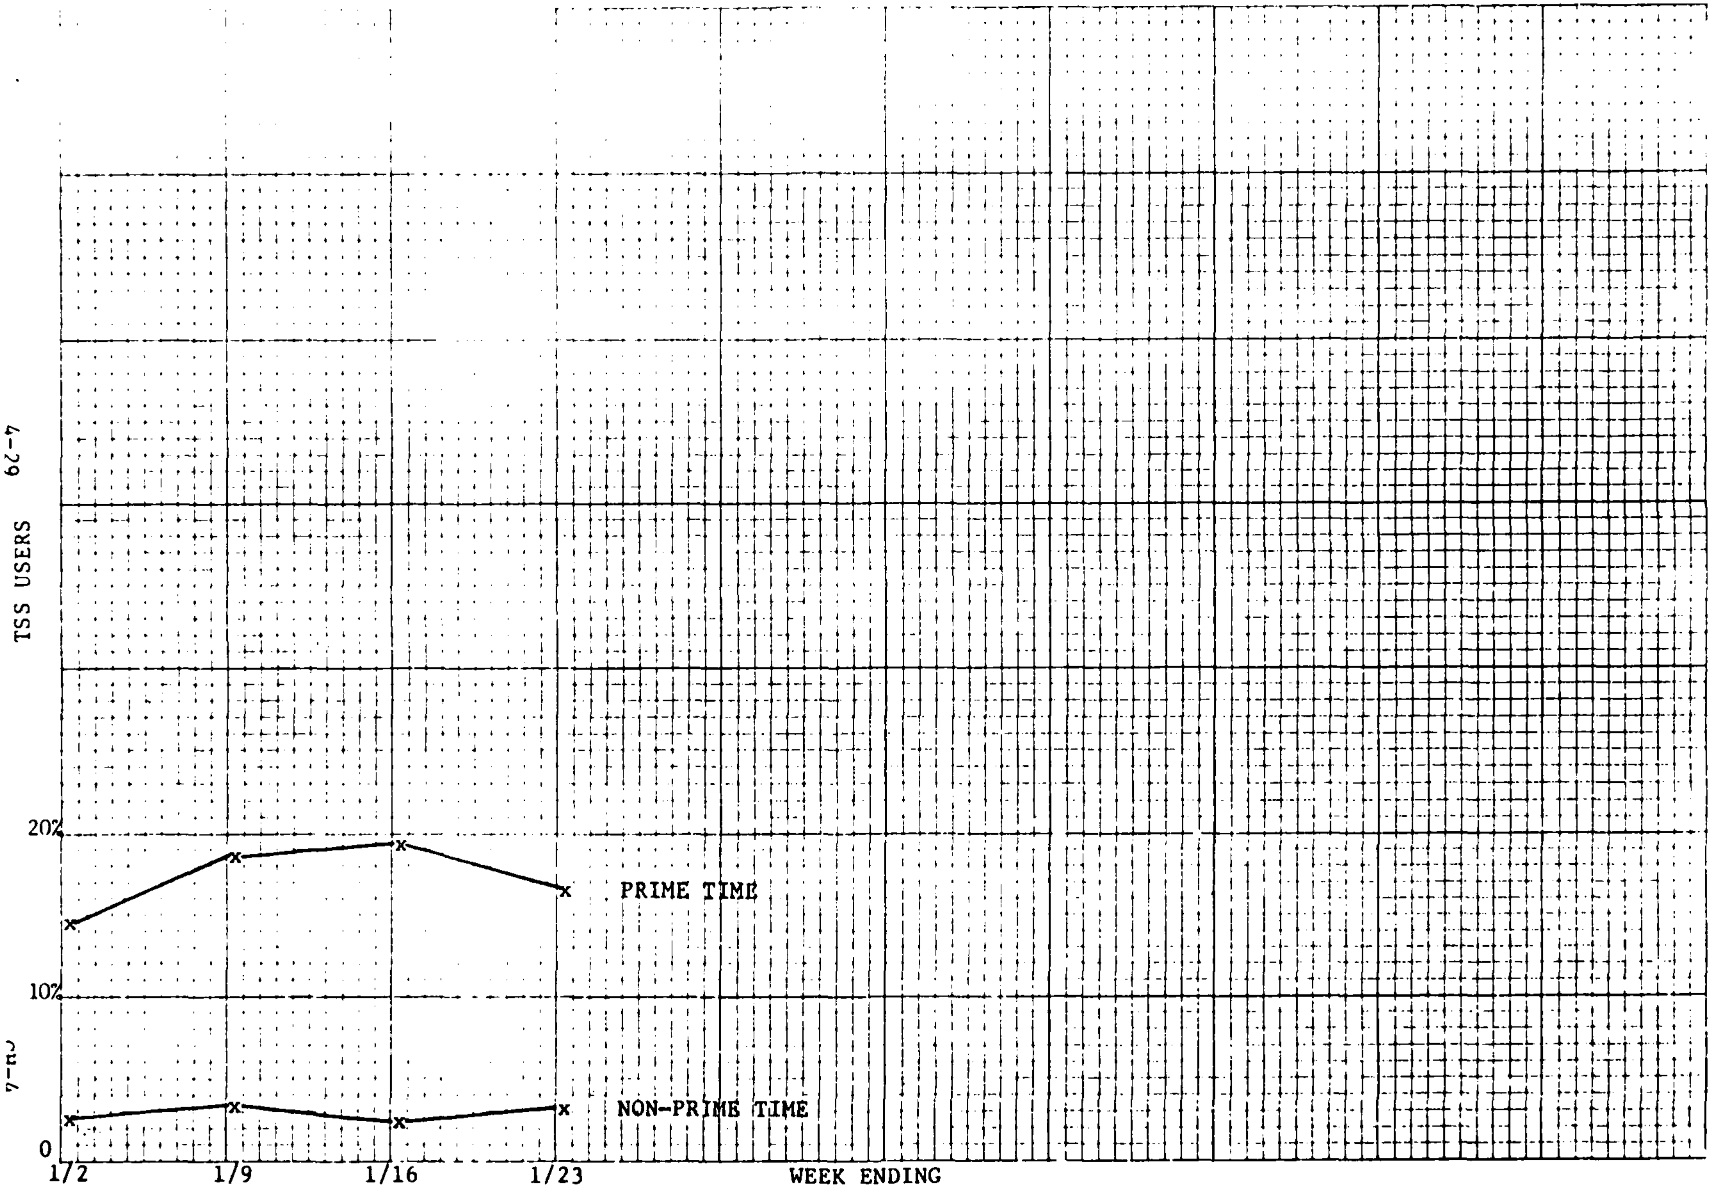

Figure 4-13. (Part 2 of 2)

THIS PAGE LEFT INTENTIONALLY BLANK

Table 5-1. Required Trace Type for Each Monitor

Monitor #	Monitor	Required Trace Type (Octal)
0	Memory Utilization Monitor (MUM)	10, 11, 46, 51, (Idle Monitor traces optional)
1	Mass Storage Monitor (MSM)	7, 15, 73*, 76*
2	CPU Monitor (CPUM)	51,63*,70*
3	Tape Monitor (TM)	50, 51, 52
4	Channel Monitor (CM)	4, 7, 15, 22 (Idle Monitor traces optional)
5	Communications Analysis Monitor (CAM)	14*, 15
6	GRTS Monitor (GRTM)	62*
7	Transaction Processing System Monitor (TPSM)	0, 1, 2, 4, 5, 6, 13, 42, 51, 65, 74*
8	Idle Monitor (IDLEM)	0, 1, 2, 3, 13, 16, 21, 22, 37, 65
A	TSS Monitor (TSSM)	74*

*These are not standard traces. They are specially created by the GMC or by an editing of the GCOS TPE Subsystem in the case of trace type 74. Trace types 63, 70, 73 and 76 are direct transfers into the GMC and as such are not required to be active via the $ TRACE card in the system boot deck. Trace types 14, 62 and 74 do use the System Trace Function and require the Trace Number to be active on the $ TRACE card.

Table 5-2. Abort Codes (Part 1 of 4)

B2 -	Illegal SNUMB on MSM data card (more than 5 characters).
B3 -	More than 5 SNUMBs for MSM SNUMB option.
BC -	Illegal request on limited connect option.
BK -	Backspace of the full data tape was bad. Multireel will not be collected. Check for tape drive problems.
BS -	Bad tape status. Check condition of tape and rerun job.
C1 -	CPU Monitor turned off but SNUMB input requested on the data card.
C2 -	Illegal SNUMB (more then five characters) on data card for CPU SNUMB option.
C3 -	More then three SNUMBs for CPU Monitor on data card.
CD -	Illegal character in CAM special option.
CE -	Console message garbled. Check console sheet and check with operator.
CM -	Cannot move out of the growth range of TSS.
CO -	CAM turned off but special option requested.
DK -	No multireel disk file was present. Use a $ FILE card in the JCL or use the M9 option to turn off multireel capability.
DR -	Disk read-in. End-of-reel processing was bad.
DS -	Bad status of the disk write.
ER -	Wrapup record could not be written.
ET -	More than two terminals requested for CAM special option.
FN -	The IOS accounting modification could not be found. Call CCTC.
GC -	No GRTS control card.
GD -	No FEP I/O can be performed.
GM -	Needed memory for GRTS Monitor denied. Increase $LIMIT card.
GO -	GRTS Monitor illegal data card.
GS -	Extra SSA is not available for GRTS Monitor. Check $ LIMIT card.

Table 5-2. (Part 2 of 4)

MO-M8,MA - Monitor was not turned off and not present on the R* file. Any monitor not contained on the R* file must be turned off via the data card option. The number following the M is the monitor that was not turned off.

MM - The dispatcher hook has already been inserted. Another version of GMC must already be in execution.

N1 - The CPU Monitor hook code could not be found. See subsection 5.2.3.

N2 - Sufficient patch space is not available in .MDISP to run the CPU Monitor. See subsection 5.2.3.

N3 - DNWW/MDNET patch location could not be found. See subsection 5.2.6.

N4 - Sufficient patch space is not available in DNWW/MDNET to run the Communications Analysis Monitor. See subsection 5.2.6.

N5 - MSM patch for SSA cache analysis not found (AOS .CRTDL). See subsection 5.2.2.

N6 - MSM patch for SSA cache analysis not found (AOS .CRTBH). See subsection 5.2.2.

N7 - MSM patch space in .MDISP not sufficient for monitoring SSA cache. See subsection 5.2.2.

N8 - CPU Monitor hook code for subdispatch could not be found. See subsection 5.2.3.

NA - CPU Monitor hook code for dispatcher queuing could not be found. See subsection 5.2.3.

NF - The Dispatcher hook code could not be found. Call CCTC/C751.

NS - A CAM abort because it could not find NSIP (# of special interrupts) address in .MDNET.

NR - A CAM abort because it could not find R01XCT (number of lines found waiting mailbox) instruction.

OE - An error in an off option was encountered. Check the data cards. There is either an illegal character on the data card or a monitor which was not compiled in the R* file (see assembly listing) has not been turned off.

Table 5-2. (Part 3 of 4)

OK - All went correctly.

OL - Overlap of disk file. Increase size of disk file. Check if operator failed to respond to tape mount message during multiprocessing.

OV - A tally overflow occurred in the MUM.T10 subroutine. Increase the size of the data area within subroutine MUM.T10, variable SIZEBUF.

RS - Routine depth requested exceeded table length.

RW - Error in initial rewind. Check tape and drive.

SB - End-of-reel processing was bad. Check tape and drive.

SD - Error setting of density.

SF - Limited reel option termed successfully.

SQ - Sequence error in the physical records.

S1 - Subroutine MUM.T10 missing

S2 - Subroutine MUM.T46 missing

S3 - Subroutine CM.T07A missing

S4 - Subroutine CPU.T70 missing

S5 - Subroutine CM.T04A missing

S6 - Subroutine CM.T22A missing

S7 - Subroutine TM.T50 missing

S8 - Subroutine CAM.T14 missing

S9 - Subroutine GRT.T62 missing

SA - Subroutine IDL.TRCS missing

SC - Subroutine IDL.T21 missing

SD - Subroutine TPE200 missing

SG - Subroutine TSS.COL missing

Table 5-2. (Part 4 of 4)

TE - The start/stop times appear improper. Check data card.

TL - Trailer record write was bad. Check tape and drive.

TS - An OK abort directed by a time to stop command.

TW - The tally word has been garbled.

TO-T8,TA - Required system trace is not on. The number following the T indicates the monitor having the problem.

Table 5-3. GMC Catalog and File Index (Part 1 of 4)

SEGMENT #	FILE	REQUIRED	FUNCTION
1	GMF.TOP	Y	Read data card, initialization, find hook in dispatcher, and create initial record
2	MUM.INIT		Initialize Memory Monitor
3	MSM.INIT		Initialize Mass Store Monitor
4	CPU.INIT		Initialize CPU Monitor
5	CAM.INIT		Initialize Communications Analysis Monitor
6	CM.INIT		Initialize Channel Monitor
7	TM.INIT		Initialize Tape Monitor
8	GRT.INIT		Initialize DN-355 Monitor
9	TP.INIT		Initialize TPE Monitor
10	IDL.INIT		Initialize Idle Monitor
10A	TSS.INIT		Initialize Timesharing Monitor
11	GMF.MID	Y	Ensure at least one active monitor
12	CAM.PAT		Preparation for patching DNWW/MDNET for Communications Analysis Monitor
13	CPU.PAT		Preparation for patching dispatcher for CPU Monitor
14	MSM.PAT		Preparation for patching dispatcher for MSM Cache Analysis
15	PATLOOK		Searches for patch space for CPUM, CAM, MSM
16	CPUDOIT		Patch the dispatcher for CPU Monitor
17	CAMDOIT		Patch DNWW for Communications Analysis Monitor
18	MSMDOIT		Patch dispatcher for MSM Cache Analysis

Table 5-3. (Part 2 of 4)

SEGMENT #	FILE	REQUIRED	FUNCTION
19	GMF.MON	Y	Insert the trace hook for GMC traces
20	CPU.REMO		Remove CPU Patches to dispatcher
21	CAM.REMO		Remove CAM patches to DNWW/MDNET
22	MSM.REMO		Remove MSM patches to dispatcher
23	GMF.BTM	Y	Remove the trace hook, do all IO control
24	IDL.TRCS		Processes traces (octal) 0,1,2,3,13,16,22,37,65 for Idle Monitor
25	IDL.T21		Processes trace (octal) 21 for Idle Monitor
26	MUM.T10		Processes traces (octal) 10,11,51 for Memory Monitor
27	MUM.T46		Processes trace (octal) 46 for Memory Monitor
28	CPU.T70		Processes traces (octal) 51,63,70 for CPU Monitor**
29	TM.T50		Processes traces (octal) 50,51,52 for Tape Monitor
30	CAM.T14		Processes traces (octal) 14 and 15 for CAM*
31	CM.T04A		Processes trace (octal) 4 for Channel Monitor
32	CM.T22A		Processes trace (octal) 22 for Channel Monitor
33	CM.T07A		Processes traces (octal) 7,15,73,76 for Channel Monitor and Mass Store Monitor**

Table 5-3. (Part 3 of 4)

SEGMENT #	FILE	REQUIRED	FUNCTION
34	GRT.T62		Processes trace (octal) 62 for GRTS Monitor*
35	GRT.COL		Interfaces with the DN-355
36	TPE200		Processes traces (octal) 0,1,2,4,5,6,13,42,51,65 and 74 for TPE Monitor*
36A	TSS.COL		Captures trace (octal) 74 for TSS Monitor*
37	RUN.GMF		JCL to run a GMC R *
38	GMF.OBJ		File to contain a GMC R *

The following set of files are a series of JCL files under the catalog B29IDPXO/GMFCOL/MAKE used to create different GMF R* monitors. The numbers in the name are the corresponding monitor number (see subsection 5.5.1):

SEGMENT #	FILE	REQUIRED	FUNCTION
39A	M02		Memory and CPU Monitors
39B	ALL		Total GMC
39C	MUM		Memory Monitor
39D	CPU		CPU Monitor
39E	TM		Tape Monitor
39F	MSM		Mass Store Monitor
39G	M0258		Memory, CPU, Communications and Idle Monitors
39H	M148		Mass Store, Channel and Idle Monitors
39I	CAM		Communications Analysis Monitor
39J	CM		Channel Monitor
39K	GRT		DATANET-355 Monitor

Table 5-3. (Part 4 of 4)

SEGMENT #	FILE	REQUIRED	FUNCTION
39L	M48		Channel and Idle Monitors
39M	M14		Mass Store and Channel Monitors
39N	M56		Communications and DATANET Monitors
390	TPE		TPE Monitor
39P	TSS		TSS Monitor
39Q	M025		Memory, CPU and Communications Monitors
39R	M01245		Memory, Mass Store, CPU, Channel and Communications Monitors

*Trace types 14,62 and 74, are not standard. They are internally generated (IT) traces.

**Trace types 63,70,73 and 76 are not standard. They are direct transfer (DT) traces.

information from the Peripheral Allocator is reported only when the Peripheral Allocator is in memory and a Memory Monitor trace is about to be generated. For this reason, not all Peripheral Allocator queue changes will be reported. In order to reduce the amount of information being collected, a job's status in the Peripheral Allocator's queue is reported only for new jobs, when a job has changed activity, or when its status has changed.

After reporting any Peripheral Allocator status information, the MUM will next report the status of every job waiting for or currently using memory. Information such as the SNUMB, IDENT, USERID, Activity Number, memory demands, current memory address, whether the job is in memory or waiting for memory, and whether the job is a system program or user program is collected. This information is reported for each job only if a change has occurred from previous information that was reported. In addition, the current amount of CPU and IO time used by a job is reported in every MUM trace that is generated.

The MUM will consider a job to be a system job if it has a program number less then octal 10, or if it has no J* file and requires privity. Since the user may want to consider other jobs to be system jobs, such as HEALS or VIDEO, the data reduction program allows the user to extend this definition of system jobs (see section 6).

5.2.2 Mass Storage Monitor. The Mass Storage Monitor (MSM) is used to collect data on usage of peripheral resources. For a detailed description of reports containing data collected by this monitor, see section 7.

When the user wants MSM to be active, it is essential that trace types (octal) 7 and 15 are enabled in the computer system boot deck on the $ TRACE card. MSM processes trace types 7, 15, 73, and 76 to build its own records which are passed to ER. A separate discussion of the format of the MSM collected data records is contained in subsection 5.4.3. As has been stated earlier, trace types 73 and 76 are direct transfer traces created by the GMC, and as such need not be defined on the $ TRACE card. The MSM requires that at least the following segment numbers from table 5-3 be used to generate the GMC R* file: 1, 3, 11, 14, 15, 18, 19, 22, 23, and 33. The complete process for generating an R* file is described in subsection 5.6. If the system being monitored by the Mass Store Monitor contains SSA Cache Core, two new direct transfer traces, are created by the Mass Store Monitor in order to collect sufficient data to be able to analyze the operation of SSA Cache Core. These traces are created only if SSA Cache Core is configured. The Mass Store Monitor searches the dispatcher for a AOS .CRTDL instruction and then inserts code to make a direct transfer into the GMF. In addition, an AOS .CRTBH instruction is also searched for so that another direct transfer into the GMC can be generated. The first instruction locates the area of the dispatcher where it has been determined that the required SSA module is not in the SSA Cache Buffer and needs to be loaded from disk. The second instruction

AD-A126 234 GENERALIZED MONITORING FACILITY USERS MANUAL CHANGE 4 (U) COMMAND AND CONTROL TECHNICAL CENTER WASHINGTON DC 01 JAN 83 CCTC-CSM-UM-246-82-CHG-4 2/3

UNCLASSIFIED F/G 9/2 NL

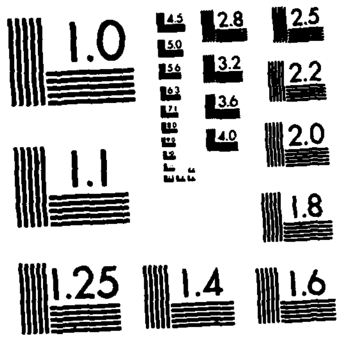

MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

locates that area of the dispatcher where it has been determined that the required SSA module is in the SSA Cache Buffer. Section 2.6.2 (MSM.PAT) completely describes the procedure for locating these instructions. If GMC cannot find these instructions between these locations, it will abort with an N5 or an N6 abort. If this problem occurs, the dispatcher code should be examined to see if this instruction has been moved or modified. If so, the code in GMC will need to be altered.

Upon finding the above sets of instructions, GMC searches the dispatcher area for 8 free locations in which to create two new direct transfer traces. Section 2.6.2 (MSM.PAT) completely describes the procedure for locating the patch area. If 8 words of free space are not found, an N7 abort will occur. In this case, the user should examine the patch deck and a listing of the patches on the total edit tape to see if a large number of patches have been made to the dispatcher. If this is the case, the dispatcher code will need to be reassembled in order to remove these patches or else the Monitor will not be able to be utilized. The user does have another alternative. This alternative involves patching word 0 of the dispatcher in order to generate a user patch area. The patch involves the setting of bit 2 to a 1 in word 0 of the dispatcher. No other modification by the user is necessary.

The MSM collects sufficient information so as to be able to completely monitor the usage of the entire disk subsystem, the usage of SSA Cache core and the usage of FMS catalog cache, when active. When either the MSM or CM is active, a record containing device names and addresses is written at the beginning of the GMC run and periodically afterward if device names change. This is done only for mass store devices. Every time the system issues a connect request to a tape drive or disk drive, sufficient information is collected so as to be able to identify who is issuing the connect, the file being connected to, the pack upon which the file is located, the parameter types for the file being connected to and the reason for the connect, i.e read, write, write verify, etc.

Whenever a MME is issued the MSM will check whether it is a system job issuing a MME GEFSYE. For purposes of this check a system job is considered to be any job with a program number less then octal 15. If the Peripheral Allocator is issuing the GEFSYE, information is collected as to the SNUMB the Peripheral Allocator is working for, and the file code that is being GEFSYE'd. If the GEFSYE type is a 2, 3, 4, 5, 8, 9, 10, 11, 18, 19, 20, 22, 23, 27, 28, or a 29 then additional information is collected so as to be able to determine the CAT/FILE string of the file being GEFSYE'd. This information will be used by the data reduction program to correlate file codes used by jobs to the actual CAT/FILE string being referenced by a job. Also, sufficient information is collected so as to be able to report how

many FILSYS connects are required in order for the system to be able to allocate and deallocate files requested by a job.

If FMS catalog cache is active or available space tables are being buffered in memory then the MSM will generate a record type octal 77 with sufficient data as to be able to monitor the effect of FMS catalog cache and available space table buffering. This record is generated once, for every 5000 connects issued by the system. This is not a physical trace that is being generated and, as such, does not need to be present on the $ TRACE card. Rather, it is merely a data record that is being written to tape and given the unique number of octal 77. The data record consists of a dump of some internal performance parameters maintained by the GCOS system within modules .MFSIO and .MASO4.

5.2.3 CPU Monitor. The CPU Monitor (CPUM) is used to collect data on CPU utilization. This monitor can be used only on a WW7.2 or commercial 4JS system. For a detailed description of reports containing data collected by this monitor, see section 11. When the user desires that the CPUM be active, GCOS trace type (octal) 51 must be enabled in the computer system boot deck on the $ TRACE card. CPUM processes trace types 51, 63 and 70 to build its records which are passed to ER. A separate discussion of the format of the CPUM collected data records is contained in subsection 5.4.4. Trace types 63 and 70 are direct transfer traces, created by the GMC, and as such, need not be defined on the $ TRACE card. The CPUM requires that at least the following segment numbers from table 5-3 be used to generate the GMC R* file: 1, 4, 11, 13, 15, 16, 19, 20, 23, and 28. The complete process for generating an R* file is described in subsection 5.6. The CPU Monitor searches the dispatcher for an ASA .SALT,5 instruction and then inserts code to generate a direct transfer trace into GMC. In order to capture subdispatch processor time, it also searches for a STQ .QTOD,4 instruction and then inserts code to make a direct transfer into GMC. In the T70 capture routine, the time increment will be negative for a regular dispatch and positive for a subdispatch. In order to monitor dispatcher queuing, the CPU Monitor will search the dispatcher for an LDA .STATE,4 instruction and then insert code to generate a direct transfer trace into GMC. Section 2.6.2 (CPU.PAT) completely describes the procedure for locating these groups of instructions.

If GMC cannot find the ASA .SALT,5 instruction, it will abort with an N1 abort; if it cannot find the STQ instruction it will abort with an N8 abort; and if it cannot find the LDA instruction, it will abort with an NA abort. If these aborts occur, the dispatcher code should be examined to determine if the instruction has been modified, moved, or patched. If so, the code in GMC will need to be modified.

Upon finding these instructions, GMC searches the dispatcher patch area(s) for either 8 or 12 free locations in which to create a direct

transfer trace into the GMC. Section 2.6.2 (CPU.PAT) completely describes the procedure for locating the patch area. If patch space is not found, an N2 abort will occur. See subsection 5.2.2 for a description of an alternate procedure in case the search for patch space fails.

The CPU Monitor tracks the CPU usage of all system programs and accumulates CPU usage of slave jobs into a single value (see subsection 5.4.4). If the user desires, he can specify up to five slave jobs for which he wants the CPU monitor to track CPU usage, just as it does for system jobs. Subsection 5:5.5. describes this user option.

The CPU Monitor can be operated in one of two modes. Under the standard default mode, the monitor will capture data for reporting on CPU utilization, as well as CPU dispatcher queuing. The user should refer to subsection 5.4.4 for a description of the data collected and section 11 for a description of the reports produced by this monitor. Under this mode of operation, the monitor will use approximately 3-4% of the available processor power. The user has the option of disabling the CPU dispatcher queuing portion of the monitor. Under this mode of operation, the monitor will only require about 1% of the available processor power, but the user will receive no reports concerning dispatcher queuing, lengths of queues or amount of time in queue. See subsection 5.5.5 for a description of this user option.

5.2.4 <u>Tape Monitor.</u> The Tape Monitor (TM) is used to collect utilization statistics on magnetic tape drive activity. A separate discussion of the format of the TM collected data records is contained in subsection 5.4.5. Reports containing data collected by this monitor are described in section 11.

When the user desires that the TM be active, GCOS trace types (octal) 50, 51, and 52 should be enabled in the computer system boot deck on the $ TRACE card. TM processes these trace types to build its records which are passed to the ER. The TM requires that at least the following segment numbers from table 5-3 be used to generate the GMC R* file: 1, 7, 11, 19, 23, and 29. The complete process for generating an R* file is described in subsection 5.6.

5.2.5 <u>Channel Monitor.</u> The Channel Monitor (CM) is used to measure I/O channel activity over tape and disk channels and contention to disk devices. A separate discussion of the format of the CM collected data records is contained in subsection 5.4.6. See section 8 for a description of reports containing data collected by this monitor.

When CM is active, it is essential that GCOS trace types (octal) 4, 7, 15, and 22 are enabled in the computer system boot deck on the $ TRACE card. CM processes these trace types to build its records,

which are passed to the ER. The CM requires that at least the following segment numbers from table 5-3 be used to generate the GMC R* file : 1, 6, 11, 19, 23, 31, 32, and 33. The complete process for generating an R* file is described in subsection 5.6.

Actually, when the CM is active, sufficient data is processed for obtaining reports not only from the Channel Monitor but also from the Mass Store Monitor. The only Mass Store Monitor data that cannot be collected would be the data needed to analyze Cache Memory. If the user also wants this data to be collected, he should create an R* file from the following segments (see table 5-3): 1, 3, 6, 11, 14, 15, 18, 19, 22, 23, 31, 32, and 33. In addition, the Mass Store Monitor must be made active. There is an additional option available with the Channel Monitor. This option allows the Channel Monitor Data Reduction Program to produce a CPU Idle/IO Active Report. This report is described in section 8. To obtain this report, the Idle Monitor must be included in the R* file. In addition, all Idle Monitor traces must be active. The following segments are required to generate the R* file: 1, 6, 10, 11, 19, 23, 24, 25, 31, 32, and 33.

5.2.6 Communications Analysis Monitor. The Communications Analysis Monitor (CAM) is used to measure machine and user response time and terminal usage. A separate discussion of the format of the CAM collected data records is contained in subsection 5.4.7. The complete process for generating an R* file is described in subsection 5.6. The output reports, containing data collected by CAM, are described in section 9. When CAM is active, it is essential that the GMC generated trace type (octal) 14 and the GCOS trace type (octal) 15 are enabled in the computer system boot deck on the $ TRACE card. CAM patches the DNWW (MDNET in W7.2) module, looking for a LDQ M.LID,3 instruction followed by an ANQ =0077777,DU instruction. When the monitor is terminated, CAM removes these patches from the system. The CAM requires that at least the following segment numbers from table 5-3 be used to generate the GMC R* file: 1, 5, 11, 12, 15, 17, 19, 21, 23, and 30.

The method used by the CAM to patch DNWW/MDNET is similar to that used by the CPUM to patch the dispatcher. Section 2.6.2 (CAM.PAT) completely describes the procedure for locating these instructions.

If CAM cannot find this instruction, GMC will abort with an N3 abort. If this problem occurs, the DNWW/MDNET code should be examined to see if this instruction has been moved or modified. If so, the code in CAM.PAT will need to be altered.

Upon finding this instruction, CAM then searches DNWW/MDNET patch area for 10 free locations in which to create a new system trace type 14. Section 2.6.2 (CAM.PAT) completely describes the procedure for locating this patch area. If no space is found by this search, an N4

abort will occur. In this case, the user should examine the patch deck to see if a large number of patches have been made to DNWW/MDNET. If this is the case, DNWW/MDNET will need to be re-edited in order to remove these patches or else the CAM will not be able to be utilized. In addition to the above patching, CAM.INIT also searches DNWW/MDNET for certain instructions. Beginning at 1400 octal locations from the entry point, and continuing for 5000 octal locations, CAM.INIT searches for a ANA=0777777,DL, followed by a CMPA instruction. If it does not find these instructions, it will abort with an NS abort. CAM.INIT is searching for a number of special interrupts processed (NSIP). In addition, CAM.INIT will also search for a ANA=077, followed by a CMPA=077 instruction beginning at the above CMPA location and continuing for 3000 octal locations. If it does not find this instruction, it will abort with an NR abort. In this section, CAM.INIT is searching for the R01XCT processing (number of lines found waiting mailbox). If a specific terminal's traffic is to be monitored (see subsection 5.5.3), the CAM will insure that no more than two terminal IDs have been included. Invalid terminal IDs, extra terminal IDs or terminal IDs without the CAM input option request will cause the GMC to abort with a CD, CO, or ET abort. See table 5-2 for the meaning of these aborts.

The CAM also uses the GCOS trace type 15 (octal) to check for any JDAC processing or any other line switching which may occur.

5.2.7 GRTS Monitor. The purpose of the GRTS Monitor (GRTM) is to collect statistical data to be used in evaluating the performance of the DATANET 355 Front-End Processor (FEP). This data includes CPU Utilization, Response Time, and Terminal Performance. The collected information is sent to the GMC within the H6000, which writes the data to tape. A separate discussion of the format of the GRTM collected data records is contained in subsection 5.4.8. This tape, containing GRTM performance data and possibly data from other monitors, is used as input to a data reduction program used to produce statistical reports (see section 10). This monitor can only be used on a WWMCCS system. It cannot be used on a commercial system.

5.2.7.1 Configuration Requirements. The GRTM will execute on H6000 system with u- to eight FEPs. The monitor is designed to run on the GRTS II Phase IIA (GRTS 1.2) FEP software.

5.2.7.2 H6000 Configuration Requirements. To run GRTM, GCOS trace type (octal) 62 must be enabled via the H6000 computer system boot deck on the $ TRACE card. The GRTM requires that at least the following segment numbers from table 5-3 be used to generate the GMC R* file: 1, 8, 11, 19, 23, 34, and 35. The complete process for generating an R* file is described in subsection 5.6.

5.2.7.3 Altering of Phase II-A Software. To use the GRTM, the user must modify the standard GRTS software. It should be noted that in Release WW7.2 the alter cards needed to modify the standard GRTS software are included within the standard release. The user must check the FMAC module to insure that variable FMON has been set to 1. The setting of this variable will ensure that the required alters are assembled into the GRTS system. The FMAC module must be recompiled and the macro library reloaded. Finally, all the GRTS modules should be recompiled.

5.2.7.4 FEP Configuration Requirements. The modified GRTS II software will produce performance data. These modifications are then enabled in the software through the use of the following control card during FEP system initialization:

CC1

```
1      8      16
$      GOPT   RCSMON
```

When the monitor is not running, the FEP will function normally. Execution of the GRTM is initiated by the H6000 as it connects to each FEP to be monitored. At that time, instructional parameters are sent to each FEP to be used in determining the amount of buffer space needed for the collection of the statistical data.

The GRTM software when configured will require an additional 1,000 (decimal) words of DATANET main memory to execute. This core requirement can grow to as much as 2,500 decimal words depending upon the input options selected. The main portion of the monitor code will be resident within the GRTS II FSUB module with additional patches being incorporated within the FCCP, FEXC, FCIP, FINT, FICM, and FHCB modules. It should be noted that when the monitor is not assembled, the 1-2K of core required for the monitor will be available for buffer space. This should be kept in mind when reviewing the output reports.

5.2.7.5 Interface Requirements. During its initialization phase, the GRTM software will attempt to log onto the H6000 system via a pseudo-terminal that is used in its interaction with the host-resident GMC program. Due to .MSECR requirements, the pseudoterminal attempting to gain access to the system must be on a .MSECR configured SYSTEM HSLA subchannel. The Physical Terminal Address (PTA) used by the pseudo terminal is unique and therefore cannot be configured in the GRTS II configuration deck for any other purpose.

The format of this PTA word is as follows:

BITS	MEANING
0-3	IOM CHANNEL NUMBER OF HSLA
4-5	HSLA NUMBER (1,2,3)
6-10	HSLA SUBCHANNEL NUMBER (0-31)
11-15	POLLED SCREEN NUMBER
16-17	MUST BE ZERO

As an example, a PTA value of 317200 would be used to identify the pseudo terminal as being on:

IOM CHANNEL NUMBER OF HSLA	= 6
HSLA NUMBER	= 1
HSLA SUBCHANNEL	= 29
POLLED SCREEN NUMBER	= 0
MUST BE ZERO	

This is the same format used to describe subchannels within .MSECR.

A user must insure the setting of the PTA value as circumstances demand. By using the format shown above, the symbolic location PTA located in the FSUB module may be altered to reflect the user's own PTA value. The FSUB module must then be reassembled prior to bootloading the DATANET.

5.2.7.6 Abort Codes. The general abort codes listed in table 5-2 apply to specific options. Abort codes created specifically for the GRTM are listed below.

GS = SSA processing error caused by missing or invalid $LIMITS card

GC = Invalid GRTM options on data card file (I*)

GC = Missing GRTM option card

GD = No FEP I/O can be performed

GM = GEMORE unsuccessful in getting needed buffers

NOTE: GM aborts occur if another program occupies continuous memory just above that of the GRTM when buffer space is requested using MME GEMORE. Increasing the $LIMITS memory value or loading the GRTM immediately after booting the system will enhance the chances of getting DATANET buffers.

5.2.7.7 DATANET Monitor Software Description. The GRTS II system operating within the DATANET will be used in the collection and recording of various internal resource information which is then sent

to the GMC executing in the host (6000) system. Monitoring functions within the DATANET have been separated into three basic monitoring categories:

a. CPU and Resource Monitoring

b. Host/FEP Response Time

c. HSLA Subchannel Monitoring

Monitor information will be passed between the DATANET and the GMC program executing in the host using the normal FICM interface.

5.2.7.7.1 DATANET-HOST Interface. The GRTM executing in the DATANET will be in the form of a pseudoterminal connected to the DATANET. Once initialized, the pseudomonitor terminal will be as any other remote terminal connected to a Direct Access (DAC) program in the host.

As part of GRTM initialization, the pseudo terminal will issue a connect-to-slave request to the DAC GMC program in the host. The DAC connect name requested will be the special monitor name of GRTMN'X' with the 'X' being a number from zero through seven which corresponds to the DATANET (0-7) that issues the connect request.

The connect-to-slave request will remain outstanding until it is answered by an inquiry issued by the GMC program. Once the request has been answered, the GRTM/GMC program connection will be made.

Since the pseudo monitor terminal will not be physically connected to the DATANET (i.e., on an HSLA S/C) the need for a special monitor "Terminal Control Block (TCB)" becomes necessary.

The special monitor TCB is necessary in order to utilize the normal FICM interface in the passing of monitor information to the GMC program.

5.2.7.7.2 Monitoring of CPU. The DATANET Monitor will collect CPU and various resource information by the placement of conditionally assembled instructions at appropriate points in the ICM and EXC modules of the GRTS II software. The information to be collected at the CPU level is independent of the HSLA subchannel and is sent to the host based upon a predefined time sampling increment for writing buffers to the host collector program. For a given buffer sent to the host, the following information is included:

a. Time Idle - The amount of time in milliseconds spent in the exec idle loop since the last buffer was sent to the host.

b. Buffer Denials - A cumulative count of buffer denials. This is a count of the number of times that the GRTS II software was unable to get a buffer for a user.

c. Buffer Availability - The number of 18-bit words currently available for buffers.

d. Number of Users - Count of the number of users currently logged on the system.

e. Number of Transactions - Two counts are maintained and sent to the host:

(1) The accumulated number of transactions sent to the host, and

(2) The accumulated number of transactions received from the host.

f. Size of Transactions - There are two counts maintained and sent to the host:

(1) A cumulative count of the number of 36-bit words sent to the host, and

(2) A cumulative count of the number 36-bit words received from the host.

g. Host RSVPs Count - A cumulative count of the number of host RSVPs received by the DATANET.

h. Buffer Requests - A cumulative count of the number of times the buffer allocation routine was called.

5.2.7.7.3 Host/DATANET Response Time. The GRTS II Monitoring of the FEP/Host Response Time will be measured on a program name basis. For this monitoring, conditional coding will be added to the FICM module to detect the various requests to the Host. Each time that either a "Connect-to-Slave" or "Disconnect" request is detected, the following formatted entry will be made into the response time buffer area:

a. Function Code

b. Type of Device/Line ID

c. Time Stamp

d. DAC Name (for connect-to-slave requests only)

The Function Code (i.e., connect-to-slave, disconnect) will identify the type of request with the line number entry identifying the logical line number of the device.

The GRTM will capture the following requests:

a. Accept Direct Input (355 asking H6000 to accept data)

b. Input Accepted (input received by the H6000)

c. Send Output (355 requesting continuation of output)

d. Output Received (355 has received data from H6000)

e. Output Started (355 has begun transmitting data to terminal)

f. Output Complete (355 has completed transmission of data to terminal)

g. Accept Direct Output (H6000 has told 355 it has data to send)

h. Accept Direct Output, Then Input (H6000 has told the 355 it has data to send and expects input)

Each time one of the above requests or responses is detected, a response time buffer entry of the following format is made:

a. Function Code

b. Type of Device/Line ID

c. Time Stamp

In placing both the line number and the time stamp in every collector buffer entry, response times between the various request cycles of each DAC program executing the host will be effectively monitored.

5.2.7.7.4 <u>Terminal Monitoring.</u> GRTS II Terminal Monitoring will be on a HSLA subchannel (S/C) basis. Every monitored HSLA S/C will be allocated a four word (18 bit) record area within the output buffer where the various monitor information for the S/C is accumulated. Each word within the S/Cs record will be a "predefined" location where the various counts for the S/C are held. The first word of each S/C record will contain information as to the HSLA number and the S/C number to which the record belongs.

The GRTM will update these various counts dynamically as they occur within the DATANET.

5.2.8 Idle Monitor. The Idle Monitor (IDLEM) is used to collect data concerning CPU activity. This monitor can only be used in conjunction with the MUM or the CM. It should not be activated if one of the two aforementioned monitors is not active. If the Idle Monitor is present on the R* file and active and if, in addition, the MUM or CM is not active, then the IDLEM will automatically be turned off. The user should read subsections 5.2.1 and 5.2.5 for information concerning the use of the IDLEM. A separate discussion of the format of the collected data records is contained in subsection 5.4.9.

This monitor generates an excessive number of traces and significant overhead. It should be activated only after the user ensures that the reports produced because of its presence are an absolute necessity for the evaluation. In most cases, this monitor should not be required.

5.2.9 Transaction Processing System Monitor. The GMC Transaction Processing System Monitor (TPSM) is used to collect data on the performance of the GCOS Transaction Processing Executive (TPE) System. A separate discussion of the format of the TPSM collected data records is contained in subsection 5.4.10. The reports containing data collected by TPSM are described in section 12.

When TPSM is active, the required traces must be enabled in the computer system boot deck on the $ TRACE card (see table 5-1). A sample of the reports and run time procedures for the data reduction program can be found in the Transaction Processor System section 12. The TPSM requires that at least the following segment numbers from table 5-3 be used to generate the GMC R* file: 1, 9, 11, 19, 23, and 36. The complete process for generating an R* file is described in subsection 5.6.

NOTE: The TPSM cannot be run concurrently with the TSSM and should only be used on a WWMCCS WW7.2 system. It should not be used on a commercial release.

5.2.9.1 TPS Trace Collection. The TPSM is unlike most other GMC monitors in that monitoring of the Transaction Processing System is controlled via the operator console. Prior to collecting data, the user must alter the TPS (see subsection 5.2.9.2) and must also create a usable GMC R* file (see subsection 5.6). Once these actions are performed and a GMC execution is started, the user must still perform one additional action before data collection can begin. The TPSM is turned on or off by the console operator via the TP MESS command. The operator must request "TP MESS". When the console responds "TP MESS?", the message "TRACE ON" is entered to start processing traces, or "TRACE OFF" to discontinue trace processing. This procedure can be repeated as often as desired. The TPSM and the TSSM are the only GMC monitors that can be turned on or off while the GMC is physically executing.

5.2.9.2 Modifying the Transaction Processing System. To use the TPSM, the user must alter the Transaction Processing System. An alter file is provided with the GMF software delivery. The file name is B29IDPXO/SOURCE/TPEALT. This file contains all alters and associated JCL necessary to modify the standard TPS. These modifications apply only to the WW7.2 version of TPS. The user will need to make minor modifications to the file so as to correctly reference any permanent files that are required.

5.2.10 Timesharing Subsystem Monitor. The Timesharing Subsystem Monitor (TSSM) is used to measure TSS performance. Section 15 details those reports available from the data collected by this monitor. This monitor should be used only on a WWMCCS WW7.2 system. If desired to be run on a commercial release, care should be exercised to ensure that all alters are located correctly (see subsection 5.2.10.1).

When TSSM is active, GCOS trace type 74 (octal) must be enabled on the boot deck $ TRACE card. The TSSM causes the trace to be taken from many points in TS1; the collector builds its records which are then passed to the ER for buffering. An example of the record format appears in subsection 5.4.11. TSSM requires the following segment numbers from table 5-3 be used to generate the GMC R* file: 1, 10a, 11, 19, 23 and 36a. Subsection 5.6 shows how to generate an R* file.

THIS PAGE LEFT INTENTIONALLY BLANK

If all TSSM reports are wanted, then both the CPU Monitor and Channel Monitor are additionally required: GCOS trace types (octal) 4, 7, 10, 11, 15, 21 and 22 must also be enabled. Restricting Channel Monitor activity to SNUMB TS1 will conserve tape and GMC overhead. The TSSM and TPSM cannot be run concurrently.

The TSS Monitor consists of two major pieces of software: a data generation program loaded into TSS and a data capture routine loaded along with GMC routines. The TSS Monitor is used to analyze periods of poor TSS response to determine which users are affected, when and for how long they are affected, and what the possible causes of the poor response are. To collect information for such analysis, probes are inserted throughout TSS to gather individual pieces of information which, when combined by a data reduction program, yield the desired result. When active, the TSS Monitor is an area of code entered via transfer instructions planted throughout TSS by an initialization routine. After the TSS Monitor is entered from a probe point, it makes an entry in the GCOS trace table and then returns to the TSS process which was executing. The places where transfers are made fall into the following categories:

- o Session identification and correlation - retrieval of line ID, DRUN, ID, USERID and User Status Table (UST) address (the unique identifier which is used to correlate traces coming from an individual session)

- o Type of user interaction - retrieval of subsystem names, indication of build mode, memory sizes, derail codes

- o Terminal I/O - a single location where I/O is started and a number of locations where courtesy calls for I/O complete are paid

- o Disk I/O to user files - a single location where I/O is started and a number of locations where courtesy calls are paid for I/O complete

- o FMS service - pairs of locations, one where TSS issues a request for service, and the other where the service is complete; the second location also gives an FMS status to tell whether the service was actually performed or why it was denied

- o Processor allocation - locations for entering the processor allocation process, placing an entry into the subdispatch ready queue, and removing an entry from the subdispatch fault queue

- o Memory allocation - retrieval of subsystem size; recording of the flow of control in the memory allocation and swap processes

- o Errors and denials - retrieval of numeric error codes, recording of individual denial events

- o Intervals of no TSS service to individual users - locations for recording line switches, DRL TASK jobs, wait disposition for batch jobs, console interaction for master user

- o Intervals and events for the TSS executive - release of processor to allow subdispatch to start, execution of MME GEWAKE when there is no work to do, processing of TSS queue entries including TRACE ON and OFF, processing of terminal log-on requests prior to UST assignment

- o Denials seen by the TSS executive - swap space denial, memory allocation denials, deadlock with other batch jobs at the remote inquiry MME GEROUT

- o Unusual conditions - places where TSS recognizes that an error has occurred, errors such as bad terminal status, inability to assign a UST within 15 seconds, maximum number of users or VIPs logged on

5.2.10.1 Locations of TSS Trace Points. The following is a list of trace numbers, modules and symbolic locations provided with the TSS Monitor source code. The monitor has these locations coded as offsets relative to the beginning of the respective modules. During execution it retrieves the TSS load map and relocates these offsets. The current monitor uses offsets for W7.2.0 level 4. In future levels and releases, the symbolic locations will probably remain constant, but some octal offsets into the modules may change. Commercial users should ensure that all offsets are correct for the release they are operating under.

Trace	Module	Location	Instruction	Description
1	TSSB	.EV7+2	LDA .TSTOD	Log-on complete

Data: UST address, line ID
Use: First correlation of UST address with line ID (UST address used in all future traces for this session)

2	TSSC	DRM998+1	TRA 0,1	Swap space denied

Data: UST address
Use: Trigger resource denial report entry

3	TSSD	PATDP+3	LREG REGSAV	PAT denial

Data: UST address, flag (=0, no PAT space; |0, duplicate file name)
Use: Trigger resource denial report entry

4	TSSD	PD621+18	LDA 2,DL	GEFSYE deallocate

Data: UST address
Use: Open interval closed by trace 5

5	TSSD	PDCC+5	CMPX1 =0404600,DU	GEFSYE deallocate complete

Data: UST address, status
Use: Trigger entry in FMS report if status is octal 404600 (SMC busy) SMC wait interval is FMS status trace, 30 or 31 (wait), FMS status trace, ..., until status not octal 404600

6	TSSG	PNTTSX+9	LDAQ .LFNAM,2	Print error message

Data: UST address, error code (binary)
Use: Trigger entry in error report, possibly conditionally based on error number

82	TSSH	COMCL+4	CMPA =040000,DU	Scan command list

Data: UST address, command text
Use: Connect command (e.g. RUN) with subsystem loaded (e.g. CDIN)

64	TSSH	INTRP1+1	CANQ .FBT23,DL	Interpret primitive

Data: UST address, current program stack tally, flag for 40 file permission, primitive number (1-11)
Use: If flag set, explain return to top of stack; trace primitives not recorded by traces 8, 9, 10, 11, and 85 (e.g. XCALL, combination of CALLP, EXEC, and BIN)

7	TSSH	POPPR1+2	CMPXO POPPR1	POPUP primitive

Data: UST address, current program stack tally, subsystem name (ASCII)
Use: Cause pointer in stack of subsystem names for user to be backed up; verify subsystem name in table against name in trace

85	TSSH	POPPR1+21	LDQ 2,AU	DRL CALLSS complete

Data: UST address, subsystem name (ASCII)
Use: Cause pointer in stack of subsystem names for user to be backed up; indicate that subsystem issued DRL CALLSS and must be swapped back in

8	TSSH	CAL10+2	LDA 2,DL	CALLP primitive

Data: UST address, program stack tally, subsystem name (ASCII)
Use: Cause pointer in stack of subsystem names for user to be pushed down and name recorded (tally used to verify depth; no pushdown for DRL T.GOTO or DRL CALLSS immediately followed by DRL RETURN)

5.3 Processing

The GMC requires one tape drive, 15K-24K words of storage, and, if the multireel option will be used, 300 links of disk storage to execute. The actual memory required will depend on the number of monitors selected for execution. The GMC will lock itself into memory, assuring that it cannot be swapped or moved during the run time. An initialization procedure will attempt to relocate GMC out of the memory growth range of Time-Sharing (TS1) and, in addition, relocate GMC to the high or low end of a quadrant of memory. Due to this feature, GMC can be started at any time of the day, without fear of causing memory fragmentation or degraded TSS response.

Immediately prior to beginning collection of data, GMC will attempt to relocate out of the TSS (TS1) memory growth range. On a system with

more than 256K of memory, the growth range of TSS (TS1) is assumed to be 180K, while on a system with less than 256K of memory the growth range is considered to be 100K. During this relocation procedure GMC might grow in size; however, the user need not be concerned, since GMC will release all unneeded memory after the relocation operation is completed. If GMC cannot relocate out of the growth range of TSS (TS1), the GMC will abort with a CM abort. The user can override this abort with a data card option, but he should do so with great care since TSS response might be severely degraded if GMC prevents TSS (TS1) from growing in size. The procedures for overriding the abort can be found in subsection 5.5.4.

After GMC succeeds in relocating out of the growth range of Time-Sharing, GMC will attempt to relocate itself to the high or low end of the memory quadrant. It will relocate as far away as it can, but it will not abort if it is unsuccessful. Therefore, slight memory fragmentation is still possible, but this should cause little if any problem to the operation of the computer.

If multiple copies of TSS are present, the current version of GMC will move out of the growth range of only TS1. Therefore, the possibility does exist that GMC will be within the growth range of other copies of TSS. The user can check this by doing a List Core Console command. If GMC is within the growth range of some copy of TSS, the user may want to abort the GMC.

5.3.1 Executive Routine. The Executive Routine (ER) controls the processing of the GMC. The ER controls all inputs, outputs, start and stop time, and all required error processing. The ER also hooks the GMC into the dispatcher.

The ER begins processing by reading the input parameters on the data card. The procedure for determining the required data card parameters for GMC operation is fully described in subsection 5.5. Using these parameters, ER determines which monitors will be active during the run. If a start time is specified, ER will remain idle until the start time is reached. During this time, it is possible that GMC will be swapped. The ER will cause the GMC to terminate at a specified stop time if the option is present. Multireel output is the standard default for GMC; however, the user can request GMC to terminate after 9ne or more reels of data have been collected. The ER controls all error aborts within GMC. Abort codes and reasons for the abort are shown in table 5-2. When ER begins processing, it will query the operator for the first tape number with the message:

    *FOR XXXXX WHAT IS THE MOUNTED TAPE'S NUMBER?

where XXXXX is the SNUMB assigned GMC. The operator must enter the reel number of the tape to be used for GMC output. If the operator

processed: 2, 3, 4, 5, 8, 9, 10, 11, 18, 19, 20, 22, 23, 27, 28, and 29. The format of the GMC record generated is as follows:

Word	Bits	Information
1	0-17	Size of record (variable)
	18-26	Not used
	27-35	Octal 15 (trace number)
2	0-35	-1 if program generating this record is not PALC. Otherwise this word contains a file code in bits 18-35.
3	0-35	-1 if program generating this record is not PALC. Otherwise it contains the SNUMB of the job for which PALC is working.
4	0-17	GEFSYE type
	18-20	Processor #
	21-26	Program #
	27-35	Activity #
5-n	0-35	Catalog file string name - not to exceed 40 words.

In order to monitor the catalog file string names of user files that are being processed by the File Transfer System (FTS), it is necessary to capture the occurrence of a MME GEMORE when issued by FTS. All MME GEFSYEs issued by FTS are ignored so as not to cause a conflict with the MME GEMOREs. An FTS GEMORE is processed only if it is a GEMORE for a permanent file. The record generated is identical to the MME GEFSYE record, except that word 2 is a minus one (-1) and word 3 is the file code being used by FTS.

5.4.3.5 FMS Cache Record. During the execution of MSM or CM a special record is written at preselected times during the monitoring session. These records are used to analyze FMS catalog cache (when configured) and also the effectiveness of the incore space tables for disk devices. This record is not generated on a WW6.4 system. The format of this GMC record is shown below.

Word	Bits	Information
1	0-17	Size of record (=15)
	18-26	Not used
	27-35	Octal 77 (special flag)
2	0-35	Number of cache hits (word -12 from entry point of .MFSIO)
3	0-35	Number of writes (word -11 from entry point of .MFSIO)
4	0-35	Number of reads (word-10 from

		entry point of .MFSIO)
5	0-35	Number of reads not in CC (octal 1511 from entry point of .MFSIO)
6	0-35	Number of non-320 word reads (octal 1512 from entry point of .MFSIO)
7	0-35	Number of skips caused by .SSTAK (octal 1513 from entry point of .MFSIO)
8	0-35	Number of cache clears (octal 1514 from entry point of .MFSIO)
9	0-35	Number of no hits (octal 1520 from entry point of .MFSIO)
10	0-35	Number of hits (octal 1521 from entry point of .MFSIO)
11	0-35	.CRSHR
12	0-35	Number of times buffer allocation attempted (word -6 from entry point of .MASO4)
13	0-35	Number of times buffer busy (word -5 from entry point of .MASO4)
14	0-35	Number of times available space table was already in memory (word -4 from entry point of .MASO4)
15	0-35	Number of times available space table was in memory but was busy (word -3 from entry point of .MASO4)

5.4.4 CPUM. The CPU Monitor processes the GMC generated event trace type 70.

5.4.4.1 Trace Type 70 - Standard. This GMC record allows six processors to be monitored and allows differentiation between TS1 executive processor time and TS1 subdispatch processor time. An activity is recognized as a copy of TSS if bit 13 in its .STAT1 word is set and if its SNUMB is TS2, TS3, or TS4 (TS1 always has program number 5). The check on the .STAT1 word eliminates possible confusion between legitimate copies of TSS and GEIN execution of spawn files or termination of DRL TASK jobs by the same names. If a system program has a SNUMB of $PACT, $MOLT, $POLT, $COLT, $SOLT, or $SALT its processor time is accumulated, along with that for program number six (test and diagnostics). The format for this GMC trace type record is shown below.

Word	Bits	Information
1	0-17	Size of record (79 or 173)
	18-26	Not used
	27-35	Trace number (octal 70)
2	0-35	Not used
3	0-35	Processor time (clock pulses) for program 1 ($CALC, Core Allocator)

4	0-35	Processor time (clock pulses) for program 2 ($PALC, Peripheral Allocator)
5	0-35	Processor time (clock pulses) for program 3 ($SYOT, SYSOUT writer)
6	0-35	Processor time (clock pulses) for program 4 ($RTIN, scheduler)
7	0-35	Processor time (clock pulses) for program 5 (TS1, TSS Executive)
8	0-35	Processor time (clock pulses) for program 6 ($TOLT, T&D executive; also includes time for special T&D SNUMBs)
9-16	0-35	Processor time (clock pulses) for programs 7-14 (decimal) (Transaction Processor, Log-on, FILSYS protection, WIN, DMTEX). In commercial releases, several of these words will contain no data since many of these programs are strictly WWMCCS related.
17	0-35	Processor time (clock pulses) for GMC
18	0-35	Processor time (clock pulses) for user programs
19	0-35	Subdispatch time (clock pulses) for program 5 (TS1)
20	0-35	Processor time (clock pulses) for TS2 executive
21	0-35	Processor time (clock pulses) for TS3 executive
22	0-35	Processor time (clock pulses) for TS4 executive
23	0-35	Subdispatch time (clock pulses) for TS2
24	0-35	Subdispatch time (clock pulses) for TS3
25	0-35	Subdispatch time (clock pulses) for TS4
26	0-35	Miscellaneous subdispatch time (clock pulses) (expansion capability for TDS, TPE II)
27-51	0-35	Number of dispatches to those programs associated with words 2-26
52	0-35	RSCR time
53-58	0-35	Idle time for processors 0-5 (.CRIDT)
59-64	0-35	Overhead time for processors 0-5 (.CROVH)
65-69	0-35	Processor time (clock pulses) for 5 specially requested SNUMBs
70-74	0-35	Name of special SNUMB in BCD
75-80	0-35	Gate loop time for each processor. Module .MFALT deducts gate loop time from the processor time charged to jobs and adds it to overhead time reported in .CROVH.

(If the user disables the monitoring of dispatcher queuing, this is the end of the CPU record. If dispatcher queuing is left enabled, then the record will contain the following additional data:)

81-144	0-35	The dispatcher queue table
145-169	0-35	Dispatcher queue time (clock pulses) for those programs associated with words 2-26. It should be noted that words 162, 166-169 will contain no data since subdispatch queuing is currently not monitored.
170-174	0-35	Dispatcher queue time (clock pulses) for 5 specially requested SNUMBs

5.4.4.2 Trace Type 63 - Initial. The CPU Monitor generates an initial record which describes the dispatcher options that are active.

Word	Bits	Information
1	0-17	Record size (8)
	18-26	Not used
	27-35	Octal 63 (trace number)
2-5	0-35	First four words of .MDISP
6	0-35	.CRPJT+2
7-9	0-35	Not used

5.4.4.3 Trace Type 63 - Activity Termination. Whenever an activity terminates, a record is generated describing the dispatcher queue time accumulated by the activity. This record is only generated when the dispatcher queuing option is left enabled.

Word	Bits	Information
1	0-17	Record size (8)
	18-26	Not used
	27-35	Octal 63 (trace number)
2	0-35	SNUMB
3	0-35	Queue time in clock pulses
4	0-35	CPU time in clock pulses
5	0-17	Program number
	18-35	Activity number
6	0-35	RSCR stop time of job
7	0-35	Start time of job taken from .CRTOD
8	0-35	Accumulated swap time of job in clock pulses
9	0-35	Stop time of job taken from .CRTOD

5.4.4.4 Trace Type 63 - Termination Record. When the GMF is terminated and the dispatcher queuing option is enabled, the following termination record is generated.

Word	Bits	Information
1	0-17	Record size (386)
	18-26	Not used
	27-35	Octal 63 (trace number)
2-65	0-35	Queue time for all currently active programs in clock pulses
66-129	0-35	CPU time for all currently active programs in clock pulses
130-193	0-35	SNUMB for all currently active programs
194-257	0-35	Start time for all currently active programs taken from .CRTOD
258-321	0-35	Activity number for all currently active programs
322-385	0-35	Cumulative swap time for all currently active programs in clock pulses
386	0-35	RSCR clock time
387	0-35	.CRTOD clock time

5.4.5 TM. The Tape Monitor processes three GCOS system traces: 50, 51, and 52 and creates its own data collection records to evaluate the effect of these events.

5.4.5.1 Trace Type 50. This GCOS system trace is generated whenever an activity goes to the core allocator and will result in the generation of a GMC trace type 50 record. The format for this GMC trace type record is shown below.

Word	Bits	Information
1	0-17	Size of record (variable)
	18-26	Not used
	27-35	Octal 50 (trace number)
2	0-29	SNUMB
	30-35	Not used
3	0-35	Time stamp
4	0-11	Activity number
	12-17	Program number
	18-29	Urgency
	30-35	Octal 50
5-N	0-35	Words 1 and 2 of SCT entry for each tape used by this program

5.4.5.2 Special Trace Type 50. The first GMC trace type 50 record is a special trace 50 to indicate the status of all tape drives when the monitor first begins. Its structure is shown below.

Word	Bits	Information
1	0-17	Size of record (variable)
	18-26	Not used
	27-35	Octal 50 (trace number)

Program stack - described by following scheme:

```
(UST+.LFILE) contains ZERO    TALLY,AFTPTR
               TALLY  TALLY   *+n,6-n
                      DUP     1,5
                      ZERO    ASCII name, address of
                              primitive
                n=depth of stack (0-5, 0=no entries)
                primitive=3 if build mode
```

Flag for build mode - set nonzero if a primitive 3 is found

Subsystem BAR and LAL - found at offset .LSIZE from current UST address (bits 0-17)

Time type - found at offset .LTCW from current UST address (bits 27-35)

Log-on time - found at offset .LTALC from current UST address

Extra buffer memory address - found at offset .LKMS from current UST address (bits 0-17)

BTOS flags - found at offset .LPQF from current UST address (bits 18-35). If bit 18 nonzero, store EBM address in bits 0-17; otherwise ignore .LKMS

Memory demand - found at offset .LSIZE from current UST address (bits 18-35)

Subsystem size for penalty - found at location .TAMIS in TSSA (bits 0-17)

Penalty factor - found at location .TALPP in TSSA (bits 27-35)

5.4.12 Special Records. During the execution of the GMC, it sometimes is necessary to generate special records that describe the occurrence of a special event. Following is a description of these special records.

5.4.12.1 Lost Data Record. If the rate of data collection does not allow GMC to dump its internal buffers to tape or if the system develops a tape malfunction, it is possible for GMC to generate a lost data condition. When this condition occurs, a special trace is

generated in the last good record recorded on tape. The next good trace recorded will be found at the beginning of the next physical record.

Word	Bits	Information
1	0-35	000777200100

5.4.12.2 Termination Record. Upon termination, GMC writes out a special termination record. The format is described in subsection 5.6.1.

5.4.12.3 End-of-Reel Flag. When the multireel option is enabled, GMC writes the following special record header at the end of each tape.

Word	Bits	Information
1	0-35	007773400100
2	0-35	Continuation reel number (BCD)

5.4.12.4 MUM Lost Data. If GMC generates a lost data condition while executing the MUM, the MUM generates a special flag in the header. The format of this flag is described below.

Word	Bits	Information
1	0-17	Record size
	18-35	Octal 200110 (lost data flag)

5.4.12.5 Reconfiguraton Record. The format of this record is described in subsection 5.4.1 under the reconfiguration discussion.

5.5 GMF User Input Parameter Options

The user must start here to define his intended use of the GMC so that he can then determine his JCL and system structure.

User control over GMC is total as a complete parameter capability is provided. Sample data parameter cards are shown in figure 5-2.

The GMC functional options are:

(1) Capability to turn off any particular monitor or combination of monitors.

(2) Specifying that collection is to halt after filling 1-9 data tapes. The default is to collect an unlimited number of tapes.

CC

1

MO M3	Turn off Monitor 0 and 3
M1	Turn off Monitor 1
M1 M9	Turn off Monitor 1 and collect only a single reel
M1 *12.36,05.00	Turn off Monitor 1, start collecting data at 12.36, and collect for 5 hours
*,03.00	All Monitors are present on the R* file and are active, collection is to start at once and continue for 3 hours
+CK	All Monitors are present on the R* file, and communication traffic is to be monitored for terminal CK
M1 M4 M93	Turn off Monitors 1 and 4, and collect maximum of three reels of data
M*	Suppress abort if GMC cannot move
#VIDEO,HEALS	All Monitors are present on the R* file, and accumulate processor time in the CPU Monitor for these SNUMBs.
MO M5 M8 X1	Turn off monitors 0, 5, and 8. Collect only tape connects with the MSM and CM.
MO M2 M5 X MS2755T,RTOS	Turn off monitors 0,2,5, read second data card, turn on MSM/CM special SNUMB option to include TSS, FTS, $PALC, 2755T and RTOS.
MO M4 MS	Turn off monitors 0,4, and collect MSM/CM traces for only TSS, FTS, $PALC.

Figure 5-2. Data Card Examples

(3) Requesting complete communication data for 1 or 2 terminal IDs.

(4) Suppressing a GMC abort if it cannot move to an acceptable location.

(5) Specifying up to three SNUMBS to be processed by the CPU Monitor.

(6) Requesting that only tape connects or mass storage connects be collected, but not both. The default is to collect both types.

(7) Declaring the start and stop times of monitoring.

(8) Requesting high density tape be collected.

(9) Specifying that the Mass Store Monitor and/or Channel Monitor are to collect data only for certain jobs.

(10) Specifying that the data card options are continuing on a new card. Without this option, only a single data card will be processed.

(11) Turn off the dispatcher queuing option for the CPU Monitor.

(12) Specifying the monitoring requirements for the GRTM.

All options are listed on a single data card, unless a dual data card is stated as being used. Each option should be separated from a previous option by the presence of a single blank column.

5.5.1 On/Off Option. This option allows the user to turn off all monitors not required for his purposes. Since the GMC default is to have all monitors turned on, unless specifically turned off, and since the TSS and TPE Monitors are incompatible, the user must have a data card and at least one of these two monitors must be turned off. The code format to turn off a given monitor is:

M0 = Memory Utilization
M1 = Mass Store Monitor
M2 = CPU Monitor
M3 = Tape Monitor
M4 = Channel Monitor
M5 = Communications Monitor
M6 = GRTS Monitor
M7 = TPE Monitor
M8 = Idle Monitor
MA = TSS Monitor
MB-MF = User Developed Monitors

(See section 13 for a discussion of user developed monitors.)

Multiple monitors are turned off by listing a series of the above monitor codes, with each code separated by a single blank column.

CAUTION: The TPE Monitor and the TSS Monitor are incompatible and cannot be active at the same time.

While it is optional to turn off a monitor, a user must turn off, on the parameter card, any monitor that is not loaded in the compiled R*. Failure to do so will result in an MO-M8,MA abort. The hexadecimal digit following the M represents the monitor that is not present on the R* file, but yet was not turned off on the parameter card. Details for creation of the R* file are given in subsection 5.6.

5.5.2 Tape Selection Option. This option allows the user to specify the number of data collector tapes to be accumulated in a job run.

M9 = One reel of tape
M91 = One reel of tape
M92 = Two reels of tape
M93 = Three reels of tape
M94 = Four reels of tape
.
.
.
M99 = Nine reels of tape

If more than nine reels of tape are desired, the user must collect an unlimited number of tapes until the GMC is aborted by either a time request or an operator console abort.

When this option is used, GMC will terminate normally with an SF abort.

5.5.3 Terminal Specification Option. This option allows the user to specify the one or two terminal ID's for which it is desired to collect total terminal data. This option may only be selected if the CAM is active and causes all data seen by the H6000 for the selected terminal to be written to the GMC tape. The following format should be used for specifying the desired terminals:

+ID = Collect data for a single terminal (replace ID with the actual terminal ID)
+ID,ID = Collect data for two terminals

CAUTION: This option can possibly collect passwords and therefore the data tape will be classified to the highest level on the system. Without this option all tapes are unclassified.

5.5.4 Move Option. This option allows the user to suppress an abort if GMC cannot relocate out of the growth range of TSS. See subsection 5.3 for an explanation of the GMC relocation procedure. The proper code for this option is shown below.

M* = suppress abort

This option should not be used if the user wants to ensure that TS1 performance will not be affected by the presence of the GMC. The present version of GMC will not move out of the way of any other copies of TSS that might be active.

5.5.5 CPU SNUMB Option. This option allows the user to specify SNUMBS for the CPU Monitor special option. The CPU Monitor will separately accumulate processor times in its data records for up to five SNUMBS. Multiple SNUMBs must be separated by a comma with no intervening blanks. The last SNUMB must be followed by a blank before any new option is requested.

#SNUMB1 = Accumulate processor time for a single job (replace SNUMB1 with actual SNUMB of a job)

#SNUMB1,SNUMB2 = Accumulate processor time for two jobs

#SNUMB1,SNUMB2,SNUMB3,... = Accumulate processor time for three, four, or five jobs

5.5.6 Connect Option. This option allows the user to suppress collection of tape connects or mass store connects within the CM and/or MSM.

?1 = Only tape connects wanted
?2 = Only mass store connects wanted

The default for this option is that all tape and mass storage connects are collected. The use of this option can significantly reduce the amount of data collected by the CM and MSM.

5.5.7 Time Option. This option allows the user to specify run time parameters. The time capability is to pre-set a time to begin data collection with a time length to run after collection starts.

*07.00,04.00 = start at 7:00 a.m., run for four hours and stop at 11:00 a.m.
*16.00 = start at 4:00 p.m., no stop specified
*,02.50 = start immediately, stop after 2 1/2 hours

Rules for this option are:

a. The time option must be the last parameter on the data parameter card as the card is read left to right and time is the last entry processed.

b. Asterisk signals GMC to process the time input option.

c. Use four characters for all times in each time entry field. Time is expressed as a 24-hour clock. All zero's must be present on the parameter card.

d. If the time option specifies a start at 0900 for a 4 hour run to 1300, and GMC is not spawned until 1000, the run will still terminate at 1300. In this case, only 3 hours of data will be collected, even though 4 hours of collection was specified.

e. The user can request the following: *22.00, 04.00. This means data collection should begin at 22:00 and continue until 02:00. The GMC will handle the problem of a clock rollover.

f. GMC allocates a tape drive as soon as it initially goes into execution. It keeps this tape drive even when it goes to sleep until told to start up. Therefore, if GMC is spawned at 0700 and told to collect data starting at 1100, the tape drive will be allocated from 0700.

g. If no time option is used, the GMC will start collecting data upon entry into the system and terminate upon a console request or tape limit request. When a time option is used, the GMC will terminate normally with a TS abort.

5.5.8 <u>Specifying High Density Tape</u>. This option is no longer required. The GMC will automatically determine the tape density and make whatever modifications required.

5.5.9 <u>Limited Mass Store Monitor/Channel Monitor Collection</u>. In order to limit the amount of data being collected by the Mass Store and/or Channel Monitors, the user can request that only data being generated from certain jobs should be collected. To use this option, the characters MS must appear on the data card. If the "S" character is immediately followed by a blank, then only data for FTS, TS1 and $PALC will be captured. If the user desires, he may request that data

from five additional SNUMBs also be captured. To use this option, the SNUMBs must appear on the data card immediately after the "S" character. No blank should be present between the "S" of the MS option and the first letter of the first SNUMB. The SNUMBs must be separated by commas (no intervening blanks) and the last SNUMB must be followed by at least one blank column before a new input option is requested.

5.5.10 Request the Next Data Card. If the user is unable to fit all the aforementioned parameters on a single data card, an additional data card can be used. The user must inform the GMC that a second data card will be present by placing an "X" on the first data card. The "X" must be the last entry on the first data card and must not be placed in the middle of an input option. A given input option should be completely described before the "X" option is used. No more than two cards can be used to describe all of the standard GMF options.

5.5.11 Disable the Dispatcher Queuing Option. In order to turn off the dispatcher queuing option of the CPU Monitor, the user should type the letter "Q" on the GMF data card. See subsection 5.2.3 for an explanation of dispatcher queuing monitoring.

5.5.12 Specifying Monitoring Requirements for the GRTM. In order to collect GRTM data, an M6 must not appear on the first data card. If the M6 is omitted from the first card, then the GRTM is active, in which case additional data cards are required. The datanet specifications must be placed on a separate data card and are not directly related with the previously described options. An "X" option must not be used to indicate that the datanet option card is present. If an M6 does not appear on the standard data cards then the GMC will expect an additional data card to follow any of the standard data cards that are present. The additional data cards are free format and indicate the datanets to be monitored, the HSLA subchannels to be monitored, and finally whether response time monitoring is to be performed. By default, only CPU resource monitoring will be conducted. As stated earlier, if the entire monitoring function is to be selected, the GRTS monitor will require approximately 2K of DATANET memory. On the other hand, if only the default option is selected, the monitor will require only 1K of DATANET memory. The required parameter categories are as follows:

Dn = FEP number (n=0 to 7)

HSLAn = High Speed Line Adapter (HSLA) number (n=1 to 3 per FEP)

SCHn = Subchannel numbers associated with each HSLA entry (n=0 to 31)

Rn = Performance response time monitoring on FEP number (n=0 to 7)

Semicolons delimit each of the categories. They also indicate that more GRTS data follows. However, the last data card should not have semicolon at the end. Commas delimit subchannel sets. A "-" specifies a range of values for subchannels.

Example:

```
D1;HSLA1;SCH0-10,14,18-30;HSLA2;SCH3-15,20-30;D0;R0
```

In this example, we will monitor DATANET #1 for CPU resources (by default), for subchannel usage on HSLA 1 subchannels 0-10, 14, 18-30, and HSLA2 subchannels 3-15 and 20-30. No response time monitoring will be performed. For DATANET 0, we will monitor CPU resources (by default), no subchannels will be monitored but response time will be monitored. In order to determine the total memory used by the monitor, the following formula should be used:

```
Memory Used = 1K (default for CPU resource monitoring)
              + 32 words * (number of HSLAs)
              + 8 words * (number of subchannels requested) +
              1K (if response monitoring selected)
```

5.5.13 General Rules of the GMC Data Parameter Card. The following are general rules to be followed in defining the data card:

a. The time option, if selected, must be the last option on the data card.

b. All input elements of the eight options should be separated by a blank character.

5.6 JCL for Creation of an Object File

5.6.1 Introduction to JCL. After the user has completed a study of the options to be specified in the parameter cards described in subsection 5.5 above, the user then must build the JCL that will create the user version of a GMC object file.

The user has optional control over the creation of a GMC object file that will serve his purposes based on functions specified on the parameter cards. This optional structure minimizes the size of the GMC monitor as only necessary code is used and provides, in addition, for easier extension to the capabilities of GMC.

The four functions to be built are initialization, system patch, remove the patch, and primary monitor collection subroutines.

5.6.2 Creation of an Object File. The GMC executive routine has been subdivided into discrete sections of code based on function. In order to generate a usable GMC object file, the individual sections of code must be merged to create a single routine. This procedure has been utilized for two reasons. First, this structure permits the easy addition of new programs, i.e., monitors. Secondly, this structure allows the simple generation of a GMC containing only that code required to capture the necessary data. If a user wants to create a GMC containing only the code necessary for the Memory Utilization Monitor and Mass Store Monitor, he can easily do so. If the user then decides he does not want to run the Mass Store Monitor, he can either recreate a GMC object file or he can turn off the Mass Store Monitor via a data card. Although this procedure sounds complicated, it minimizes the size of the GMC.

When GMF is restored to the user's system, the GMC data collector program is in the form of a catalog structure. This structure is shown in figure 5-3. The GMC catalog structure as it relates to the creation of GMC R* file is shown in figure 5-4. For a description of each file function, refer to table 5-3.

As illustrated in figure 5-1, the GMC consists of an Executive Routine which initializes all the required programs, installs any required system patches, records the system patches, and removes all patches from the system when it is finished. Table 5-3 gives a breakdown of all required GMC files and any optional files. The user selects the programs he wants to monitor in the system, assembles those programs *in numerical order* (as listed in table 5-3), and runs a file-edit job to create a GMC object file. Figure 5-5 is an example of a GMC containing the Memory Utilization, Idle, and CPU Monitors. Figure 5-6 is an example of a GMC containing the Mass Store Monitor, Channel Monitor, and Idle Monitor. A $ GMAP card is required before the SELECTA for /GMF.TOP, and before every SELECTA after /GMF.BTM. Under normal operating conditions, the FILEDIT activity will contain multiple "Inconsistent Deck Name" error messages, which can be ignored. If a user should improperly create an R* file by omitting a required file, the GMC execution will abort with an S1-S9 abort, or an SA, SC, or SD abort, depending upon which routine is missing. The user should refer to table 5-2 for an explanation of these aborts.

Within the GMC release, precanned JCL has already been provided for the generation of 18 different monitor combinations. See table 5-3 for a list of the files containing this set of JCL.

The GMF is designed to be run on a HIS 6000 computer system, running with WWMCCS GCOS release 6.4 or 7.2. These releases are equivalent to the HIS commercial 2H or 4JS1 GCOS releases. The CPU, TSS and TPE Monitors must be run under a WW7.2 release.

B29IDPXO

GMFCOL

MAKE	GMF.OBJ	PATLOOK	GMF	MUM	MSM	CM	CPU	TM
ALL			GMF.TOP	MUM.INIT	MSM.INIT	CM.INIT	CPU.INIT	TM.INIT
MUM			GMF.BTM	MUM.T10	MSM.PAT	CM.T04A	CPU.PAT	TM.T50
CPU			GMF.MID	MUM.T46	MSM.REMO	CM.T22A	CPU.T70	
TM			GMF.MON		MSMDOIT	CM.T07A	CPU.REMO	
MSM							CPUDOIT	
CAM								
CM								
GRT								
TPE								
TSS								
M02								
M0258								
M148								
M48								
M14								
M56								
M025								
M01245								

Figure 5-3. GMC Catalog File Structure (Part 1 of 2)

B29IDPXO

GMFCOL (Continued)

CAM	GRT	IDLE	TPE	TSS	RUN.GMF
CAM.INIT	GRT.INIT	IDL.INIT	TP.INIT	TSS.PAT	
CAM.REMO	GRT.T62	IDL.TRCS	TPE200	TSS.COL	
CAM.PAT	GRT.COL	IDL.T21		TSSMON	
CAM.T14				TSS.INIT	
CAMDOIT				TSGMF	

Figure 5-3. (Part 2 of 2)

$:SELECTA:B29IDPXO/GMFCOL/(see below for FILE requirements and options)

REQUIRED	OPTIONAL, WHEN A MONITOR IS USED ALL PROGRAM ELEMENTS OF THE MONITOR MUST BE USED									
GMF.	MUM.	MSM.	CM.	CPU.	TM.	CAM.	GRT.	IDLE	TP	TSS
GMF.TOP	MUM.INIT	MSM.INT	CM.INIT	CPU.INIT	TM.INIT	CAM.INIT	GRT.INIT	IDL.INIT	TP.INIT	TSS.INIT
GMF.MID GMF.PATLOOK*1		MSM.PAT MSMDOIT		CPU.PAT CPUDOIT		CAM.PAT CAMDOIT				
GMF.MON		MSM.REMO		CPU.REMO		CAM.REMO				
GMF.BTM	MUM.T10 MUM.T46	CM.T07A*2	CM.T04A CM.T22A CM.T07A	CPU.T70	TM.T50	CAM.T14	GRT.T62 GRT.COL	IDL.TRCS IDL.T21	TPE200	TSS.COL

NOTE-1 GMF.PATLOOK is required only with MSM, CPU and CAM.
NOTE-2 CM.T07A is required with MSM. However, if using both MSM. and CM., then only use only one copy of CM.T07A.

Figure 5-4. GMC JCL Structure

```
$       FILEDIT SOURCE,OBJECT, INITIALIZE
$       PRMFL   R*,W,S,B29IDPXO/GMFCOL/GMF.OBJ
$       FILE    K*,NULL
$       DATA    *C,COPY
$       LOWLOAD
$       OPTION  ERCNT/500/
$       GMAP    NDECK
$       SELECTA B29DIPXO/GMFCOL/GMF/GMF.TOP
$       SELECTA B29IDPXO/GMFCOL/MUM/MUM.INIT
$       SELECTA B29IDPXO/GMFCOL/CPU/CPU.INIT
$       SELECTA B29IDPXO/GMFCOL/IDLE/IDL.INIT
$       SELECTA B29IDPXO/GMFCOL/GMF/GMF.MID
$       SELECTA B29IDPXO/GMFCOL/CPU/CPU.PAT
$       SELECTA B29IDPXO/GMFCOL/PATLOOK
$       SELECTA B29IDPXO/GMFCOL/CPU/CPUDOIT
$       SELECTA B29IDPXO/GMFCOL/GMF/GMF.MON
$       SELECTA B29IDPXO/GMFCOL/CPU/CPU.REMO
$       SELECTA B29IDPXO/GMFCOL/GMF/GMF.BTM
$       GMAP    NDECK
$       SELECTA B29IDPXO/GMFCOL/IDLE/IDL.TRCS
$       GMAP    NDECK
$       SELECTA B29IDPXO//GMFCOL/IDLE/IDL.T21
$       GMAP    NDECK
$       SELECTA B29IDPXO/GMFCOL/MUM/MUM.T10
$       GMAP    NDECK
$       SELECTA B29IDPXO/GMFCOL/MUM/MUM.T46
$       GMAP    NDECK
$       SELECTA B29IDPXO/GMFCOL/CPU/CPU.T70
$       EXECUTE
$       ENDEDIT
$       ENDCOPY
```

Figure 5-5. Creation of R* File for Memory, Idle, and CPU Monitor

```
$       FILEDIT     SOURCE,OBJECT,INITIALIZE
$       PRMFL       R*,W,S,B29IDPXO/GMFCOL/GMF.OBJ
$       FILE        K*,NULL
$       DATA        *C,COPY
$       LOWLOAD
$       OPTION      ERCNT/500/
$       GMAP        NDECK
$       SELECTA     B29IDPXO/GMFCOL/GMF/GMF.TOP
$       SELECTA     B29IDPXO/GMFCOL/MSM/MSM.INIT
$       SELECTA     B29IDPXO/GMFCOL/CM/CM.INIT
$       SELECTA     B29IDPXO/GMFCOL/IDLE/IDL.INIT
$       SELECTA     B29IDPXO/GMFCOL/GMF/GMF.MID
$       SELECTA     B29IDPXO/GMFCOL/MSM/MSM.PAT
$       SELECTA     B29IDPXO/GMFCOL/PATLOOK
$       SELECTA     B29IDPXO/GMFCOL/MSM/MSMDOIT
$       SELECTA     B29IDPXO/GMFCOL/GMF/GMF.MON
$       SELECTA     B29IDPXO/GMFCOL/MSM/MSM.REMO
$       SELECTA     B29IDPXO/GMFCOL/GMF/GMF.BTM
$       GMAP        NDECK
$       SELECTA     B29IDPXO/GMFCOL/CM/CM.TO4A
$       GMAP        NDECK
$       SELECTA     B29IDPXO/GMFCOL/CM/CM.T22A
$       GMAP        NDECK
$       SELECTA     B29IDPXO/GMFCOL/CM/CM.TO7A
$       GMAP        NDECK
$       SELECTA     B29IDPXO/GMFCOL/IDLE/IDL.TRCS
$       GMAP        NDECK
$       SELECTA     B29IDPXO/GMFCOL/IDLE/IDL.T21
$       EXECUTE
$       ENDEDIT
$       ENDCOPY
```

Figure 5-6. Creation of R* File for Mass Store Monitor, Idle and Channel Monitor

When GMF is used on WWMCCS release 7.2, or commercial release 4JS1-4JS3, the user must ensure that the value for variable "SYS64" is set to 0. This variable is defined in an "EQU" statement located in the following files: B29IDPXO/GMFCOL/GMF/GMF.TOP, B29IDPXO/GMFCOL/CM/CM.TO7A, and B29IDPXO/GMFCOL/MUM/MUM.T10. See subsection 2.6.2 for a discussion of other GMC modifications that are required under different GCOS software releases.

Having created a GMF object file, no other system modifications are required unless the GRT Monitor, the TPE Monitor or the TSS Monitor are desired. Each of these monitors require system modifications to be made prior to their use. The system modifications required by the GRTS Monitor are described in subsection 5.2.7 of this section. The system modifications required by the TPE Monitor are described in subsection 5.2.9 of this section. The system modifications required by the TSS Monitor are described in subsection 5.2.10 of this section.

5.7 JCL for Executing the GMC

The JCL needed to execute the GMC is shown in figure 5-7. The size to be placed on the $ LIMITS card depends on the number of monitors present in the R* file. The size should range from 15K to 24K, depending on the number of monitors. The load map produced on the compilation listing from the General Loader will specify the actual memory size required to load GMC. In figure 5-7, parameter card following $ DATA I* demonstrates a turnoff for the MUM, CPUM, TM, GRTM, IDLE, TPEM monitors, and a collection of an unlimited number of tapes. If the GRTS monitor is active, it may dynamically grow GMC during its initialization procedure. Because of this feature, the GRTS monitor requires a $ LIMITS card with a large SYSOUT request to force the system to obtain an extra 1K memory. Figure 5-8 shows the JCL needed to execute a GMC session which includes the GRTM. Since the GMC runs in master mode, it requires a PRIVITY card, and the operator must grant it permission to run. The DK file is optional but must be present if more than one reel will be collected. The size of the DK file is approximately 300 links but may require more than 300 links on a very busy system. This file is used to collect data during rewind of a completed data tape. When GMF loads, many load error messages will be produced. These error messages may all be ignored.

```
$       OPTION  ERCNT/999/
$       LOWLOAD
$       EXECUTE DUMP
$       PRMFL   R*,R,S,B29IDPXO/GMFCOL/GMF.OBJ
$       LIMITS  99,16K
$       PRIVITY
$       TAPE    OT,X2D,,99999
$       FILE    DK,X1R,300R
$       DATA    I*
MO M2 M3 M6 M7 M8 (optional)
$       ENDJOB
```

*NOTE - The execution report from GMC will contain multiple "nonfatal error" messages which can be ignored.

**NOTE - If high density tape collection is desired, the tape card should be modified with the addition of four commas after the tape number followed by the words DEN16:

```
$     TAPE     OT,X2D,,12345,,,,DEN16
```

and also the option M; should be on the data card.

Figure 5-7. JCL for Executing the GMC

```
$       LOWLOAD
$       OPTION      ERCNT/500/
$       EXECUTE     DUMP
$       PRMFL       R*,R,S,B29IDPXO/GMFCOL/GMF.CBJ
$       LIMITS      99,15K,,99999
$       PRIVITY
$       TAPE        OT,X2D,,,,GRTSII-DATA
$       FILE        DK,X1D,300R
$       DATA        I*
MO M1 M2 M3 M4 M5 M7 M8 M9 MA
 DO;HSLA1;SCHO-8,14,18,21-23;D1;RO
$       ENDJOB
```

Figure 5-8. GRTM JCL

Table 6-1. (Part 2 of 4)

ID Number	Histogram Title
21	Number of Activities Waiting Memory When a Processor Went Idle *
22	Memory Available When a Processor Went Idle *
23	Delay Time in the System Scheduler
24	Delay Time Until Core Allocation
25	Percent of Assigned Memory Used (Time-Corrected) *
29	Number of User Activities Waiting Memory in the Allocator Queue *
30	Number of User Activites in Memory *
31	Elapsed Time of a Busy State Processor 0
32	Elapsed Time of a Busy State Processor 1
33	Elapsed Time of a Busy State Processor 2
34	Elapsed Time of a Busy State Processor 3
35	CP Time Per User Activity
36	I/O Time Per User Activity
40	Number of Activities Waiting Memory (Time-Corrected) *
41	Number of Activities in Memory (Time-Corrected) *
42	Memory Available (Time-Corrected) *
43	Number of User Activities Waiting Memory (Time-Corrected) *
44	Number of User Activities in Memory (Time-Corrected) *
45	Total Demand Outstanding (Time-Corrected) *

* Jobs with 0 urgency are not included.

Table 6-1. (Part 3 of 4)

ID Number	Histogram Title
46	Number of Extra Activites That Could Fit in Memory Without Compaction
48	Length of an Idle State (All Processors)
49	Length of an Idle State Processor 0
50	Length of an Idle State Processor 1
51	Length of an Idle State Processor 2
52	Length of an Idle State Processor 3
53	Number of Times System Activity Swapped
54	Elapsed Time a System Activity was Swapped
55	Elapsed Time of a Busy State Processor 4
56	Elapsed Time of a Busy State Processor 5
57	Length of an Idle State Processor 4
58	Length of an Idle State Processor 5
ID Number/Name	**Plot Title**
26/PLOT1	Availability of Memory vs. Outstanding Demand In Core Allocator Queue vs. Outstanding Demand in Peripheral Allocator Queue Plus Outstanding Demand in Core Allocator Queue
27/PLOT2	Memory Shortfall in Core Allocator Queue vs. Memory Shortfall in Core Allocator Queue Plus Memory Shortfall in Peripheral Allocator Queue
28/PLOT3	Number of Activities in Core Allocator Queue vs. Number of Activities in Peripheral Allocator Queue
59/PLOT4	Average size of TSS, FTS and NCP. In commercial releases, only the TSS will be plotted.

Table 6-1. (Part 4 of 4)

ID Number/Name	Other Reports
37/PALC	Peripheral Allocator Report
38/ACTIVE	Activity Report/Excessive Resource Report/Abort Report/IDENT Report
39/MAP	Memory Map
47/OUT	Out of Core Report
---	Special Job Memory Reports
---	System Program Usage Report
---	Memory Statistics Report
---	Distribution of Urgency Over Time Report
---	Zero Urgency Job Report

cover a larger range of values. This change could be made via data cards and would not increase the size of the program.

The second method would involve increasing the size of the histogram by altering the value of TABSIZ. As long as the size requested does not exceed 50, this change can also be done via a data card. However, if an individual histogram needs to be larger than 50 buckets, the user will need to change the value of MXTBSZ. This change will require a change to source code, a recompile, and probably an increase in program size. All references to MXTBSZ must be altered. This would need to be done in the EDIT subsystem of Timesharing.

The remaining items that can be modified are the title and the vertical axis headers. The method for altering the histogram parameters is detailed in subsection 6.1.6. Table 6-2 shows the default values for all histograms.

6.1.4 Plot Options. There are three characteristics directly available to the user for each individual plot axis used.

The first characteristic, MAXNUM, is the maximum number of entries to be plotted on each vertical plot axis.

The second characteristic, YMAX, defines the upper limit of the horizontal display axis.

The third characteristic, YMIN, defines the lower limit of the horizontal display axis. The method for altering these values is explained in subsection 6.1.7. Table 6-3 shows the default values for all plots.

6.1.5 Default Option Alteration. The general format for an option request is as follows: the first card contains an action code describing the action to be taken. Subsequent cards modify report parameters for some of the action codes. All input cards are free format with the only requirements being that at least one blank space separates multiple input parameters. The very last input card must have the word "END" typed on it. This card must be present whether or not any other input options are selected. Available actions with their (default) implications are shown in table 6-4. There is no order required for the options. In reading the following sections it should be remembered that the first card for any input option must be the action code specification with no other data present on the card.

The user should take special note that if this software is executed under a WW6.4/2H GCOS release, an additional data card is required. The data card should contain the letters RN.

Table 6-2. Default Values for Histograms (Part 1 of 2)

ID #	Low Value	Interval Size	Number of Buckets
1	4	4	50
2	0	50	50
3	0	250	50
4	0	1	50
5	4	4	50
6	0	1	50
7	0	5	50
8	0	200	50
9	0	200	50
10	.95	.1	50
11	4	4	50
12	0	10	50
13	0	1	50
14	0	1	50
15	0	1	50
16	0.0	5.0	50
17	0	10	50
18	4	10	50
19	4	20	50
20	0	1	50
21	0	1	50
22	4	8	50
23	0	25	50
24	0	25	50
25	50	2	50
29	0	1	50
30	0	1	50
31	0.0	5.0	50
32	0.0	5.0	50
33	0.0	5.0	50
34	0.0	5.0	50
35	5	5	50
36	5	5	50
40	0	1	50
41	0	1	50
42	0	10	50
43	0	1	50
44	0	1	50
45	0	10	50
46	0	1	50
48	0.0	5.0	50
49	0.0	5.0	50
50	0.0	5.0	50

Table 6-2. Default Values for Histograms (Part 2 of 2)

ID #	Low Value	Interval Size	Number of Buckets
51	0.0	5.0	50
52	0.0	5.0	50
53	0	1	50
54	0	250	50
55	0.0	5.0	50
56	0.0	5.0	50
57	0.0	5.0	50
58	0.0	5.0	50

Table 6-3. Default Values for Plots

ID #	Max Size of Plot	Lower Plot Limit	Upper Plot Limit
26	Unlimited	0.	456.
27	Unlimited	0.	456.
28	Unlimited	0.	114.
59	Unlimited	0.	228.

Table 6-4. Available Report Actions and Their (Default) Values (Part 1 of 2)

HISTG - Modify a histogram (see table 6-2)

PLOT - Modify a Plot (see table 6-3)

ON - Turn a specific report on - (all reports on except Memory Map and Out of Core Report)

OFF - Turn a specific report off - (all reports on except Memory Map and Out of Core Report)

TIME - Set a timespan(s) for reporting - (total time reported)

ALLOFF - Turn all reports off except those specified - (all reports on except Memory Map and Out of Core Report)

ALLON - Turn all reports on except those specified - (all reports on except Memory Map and Out of Core Report)

ERROR - Do not stop on an option request error - (stop on an input error)

DEBUG - Program debug requested - (no debug)

ALLOC - Stop program after a specified number of memory allocations have been requested - (entire tape processed)

NREC - Stop program after a specified number of tape records have been processed - (entire tape processed)

NOUSER - Do not print USERID on any report - (USERID printed on certain reports)

IDLE - Turn off all Idle Monitor reports - (all IDLE reports on)

WASTED,CORE,IO,CPU,RATIO,URG - Changes parameters used in the Excessive Resource Usage Report - (20K,50K,30MIN, 30MIN,5,40)

ABORT - SNUMBs not to report in the ABORT Report - (all SNUMBs that abort are reported)

PLTINT - Change Interval at which plots are printed - (10 MIN)

FSTSLV - Change the lowest allowable user program number - (14 decimal)

Table 6-4. (Part 2 of 2)

MASTER - Define SNUMBs that are considered to be SYSTEM programs - (all programs with a program number less than FSTSLV). In addition, the programs $TRAX, $HEALS and VIDEO will be considered system programs.

PALC - Change the print control for the PALC report (600 secs)

END - Required as last card of input. It must be present.

SPECL - Produce the Special Job Memory Reports

RN - Processing on a WW6.4 system

MAPART - Produce a memory map only when the number of jobs waiting memory surpasses a predetermined value.

6.1.6 Histogram Alterations (Action Code HISTG). A complete description of histogram default values and their meanings is provided in section 6.1.3. In order to change any histogram parameter, the user is required to supply a series of four data cards. The first card contains the action code HISTG. The second card specifies which histogram the user wants to alter. This specification is made by inputting the histogram ID number as obtained from table 6-1. The third data card describes the parameter to be changed and the fourth card provides the new value for the parameter. The following options are available:

Card #3	Card #4
LOWVAL	A new low value
SIZE	A new maximum histogram size
BUCKET	A new bucket size
HEADER	A two-word header separated by at least one blank. Each header word must not exceed six characters in length. If one of the headers is to be blank, the word BLANK must be typed on the data card.

Two additional points must be stressed.

- o If the user wants to change multiple parameters for a single histogram, then cards 3 and 4 should be repeated as often as required. When alterations for a given histogram are completed, the user must supply a new action request. If the user desires to change another histogram, then the entire sequence of four data cards must be repeated.

- o When inputting the new parameter values, the user must consult table 6-2 in order to determine whether the parameters must be inputted as integer (no decimal point), or as real numbers (decimal point must appear on data card).

Figure 6-1 shows a standard histogram format. Column five of this histogram has the words "NUMBER" and "WAIT." These are called the header labels and are used to describe the function being reported by this histogram. It is these two words that the user may modify with the HEADER parameter card. Figure 6-2 shows the input option formats for this action code.

6.1.7 Plot Alterations (Action Code PLOT). Modifications to a plot allow the user to specify a new plot size, a new maximum horizontal axis limit and a new minimum horizontal axis limit. The default values for the existing plots are described in section 6.1.4. As with the histogram option, the user is required to supply a series of four data cards for each parameter change desired. The first data card contains the action

```
                    DISTRIBUTION COLLECTED ON SYSTEM NMCC2  AT 12:39:46 ON 80-12-15

                    THE # OF USER ACTIVITIES WAITING MEMORY IN ALLOCATORS QUEUE

INDIV.  CUMUL.   CUMUL.   INDIV.  NUMBER         PERCENT OF OCCURRENCE                 REPORT
NUMBER  NUMBER    PRC      PRC     WAIT    00   05   10   15   20   25   30   35   40   45   50      10
                                            I....I....I....I....I....I....I....I....I....I....I
 1456    1456    6.220    6.220    0-   0   IXXXXXX
 2257    3713   15.861    9.641    1-   1   IXXXXXXXXXX
 3949    7662   32.730   16.869    2-   2   IXXXXXXXXXXXXXXXXX
 5456   13118   56.036   23.306    3-   3   IXXXXXXXXXXXXXXXXXXXXXXX
 4053   17171   73.349   17.313    4-   4   IXXXXXXXXXXXXXXXXX
 2924   20095   85.839   12.490    5-   5   IXXXXXXXXXXXX
 1836   21931   93.682    7.843    6-   6   IXXXXXXXX
  878   22809   97.433    3.751    7-   7   IXXXX
  317   23126   98.787    1.354    8-   8   IX
  114   23240   99.274    0.487    9-   9   I
  133   23373   99.842    0.568   10-  10   IX
   34   23407   99.987    0.145   11-  11   I
    3   23410  100.000    0.013   12-  12   I

23410 ENTRIES TOTAL     AVERAGE =   3.40961     VARIANCE =   3.821     STANDARD DEVIATION =   1.955
```

Figure 6-1. Standard Histogram

Card	1	A
	2	B
	3	C
	4	D

where

A = The word HISTG
B = Histogram ID number (table 6-1)
C = Parameter control word
D = New parameter value. If parameter control word was HEADER, then this card must contain two words separated by at least one blank. Each word cannot exceed six characters in length. If the user desires one of the words to contain blanks, he must type the word BLANK on the card.

* Cards 3 and 4 are repeated for as many parameter changes as desired for the single histogram defined by card 2.

Figure 6-2. HISTG Action Code Format

code PLOT. The second card specifies which plot the user wants to alter. This specification is made by inputting the plot ID number from table 6-1. The third data card describes the parameter to be changed and the fourth data card provides the new value for the parameter. The following options are available:

Card #3	Card #4
SIZE	A new maximum plot size
LOWVAL	A new low value
HIVAL	A new high value

Several points must be stressed.

- o When inputting parameter values for LOWVAL and HIVAL, these values must be inputted as real numbers (decimal point must appear on data card). If a new LOWVAL or HIVAL value is selected, the user should try to ensure that (HIVAL-LOWVAL) is divisible by 114.

- o When inputting parameter values for SIZE, the value must be inputted as an integer. If the size is to be unlimited, then a -1 must be inputted.

Figure 6-3 shows a standard plot format. The maximum and minimum values of the plot are used to determine the value of each dash across the horizontal axis. In this figure, each dash has a delta value of 1.0. Figure 6-4 shows the format for this input option.

6.1.8 Turn a Report On (Action Code ON). This card allows a user to turn on a single report that is off by default. Card 1 contains the word ON and card 2 contains the NAME of the report to be turned on (table 6-1). No change to the default parameters, of the report, will be made. This option may be used only for Plots and Other Reports and cannot be used to control individual histograms. Note that only those reports and/or plots that have a specific NAME can be controlled with this option.

6.1.9 Turn a Report Off (Action Code OFF). This card allows a user to turn off a single report that is on by default. Card 1 contains the word OFF and card 2 contains the NAME of the report to be turned off (table 6-1). No change to the default parameters, of the report, will be made. This option may be used only for Plots and Other Reports and cannot be used to control individual histograms. Note that only those reports and/or plots that have a specific NAME can be controlled with this option.

6.1.10 Set a Timespan of Measurement (Action Code TIME). The timespan of data collection can cover many hours of which only a few may be of interest. This option allows a user to specify the timespan (or spans) to be displayed in all reports. For example, the user may specify that he wants to collect data from 0500 to 2200 and wants to display data only

from 0900 to 1700 in all reports. As another option, the user may request to see the memory map from 0900 to 1000, plot #1 from 1200 to 1500 and all other reports from 08[illegible] to 1700.

The user must [illegible]cify the report NAME to be affected by the time request (table 6-1). If the entire reduction (all repo[illegible]ts) are to be time controlled, a report NAME of "TOTAL" must be used. Histogram reports cannot be individually time spanned. Note that only those reports and/or plots that have a specific NAME can be controlled with this option. All time spans of plots or other reports will be bound by the total report timespan, if one is to be used. Up to five timespans for each report or plot may be specified, and they must be serially ordered. All times are expressed as SIMI where SI is the hour and MI is the minute. All time

```
                         DISTRIBUTION COLLECTED ON SYSTEM NMCC2 AT 12:39:46 ON 80-12-5

                           # OF ACTIVITIES IN CALC QUEUE VS # IN PALC QUEUE

DELTA =   1.0       A=NCALC       B=NPALC

TIME OF DAY                                              NUMBER WAITING
    0.                       28.0                        57.0

                            ^                            ^
Y---------Y---------Y---------Y---------Y---------Y---------Y---------Y---------Y---------

2                                                                                          12:49:43.00
2                                                                                          12:59:44.00
2                                                                                          13:09:44.00
2                                                                                              . .
2                                                                                              . .
2                                                                                              . .
2                                                                                              . .
2                                                                                              . .
2 A                                                                                            . .
B   A                                                                                          . .
B   A                                                                                          . .
B     A                                                                                        . .
B       A                                                                                      . .
B                 A                                                                            . .
B         A                                                                                    . .
B         A                                                                                    . .
B     B       A                                                                                . .
I       B  A                                                                                   . .
I     B     A                                                                                  . .
I     B     A                                                                                  . .
I         B A                                                                                  . .
I     B   A                                                                                16:19:42.00
```

Figure 6-3. Standard Plot

Card	1	A
	2	B
	3	C
	4	D

A = The word PLOT
B = Plot ID number (table 6-1)
C = Parameter control word
D = New parameter value

* Cards 3 and 4 are repeated for as many parameter changes as desired for the single plot defined by card 2.

Figure 6-4. PLOT Action Code Format

must be expressed as four character fields with no intervening blanks. Time is based on a 24-hour clock. If a user wants to request the time 4:07, he must input 0407. All times must include four characters.

If a start time, but no stop time, is desired, no characters should be entered after the minutes of the start time. If a stop time is requested, there must be a start time corresponding to it. If the user wants to start at the beginning of data collection and stop at some specified time, but is not sure of the start time, a start time of 0001 should be used. Figure 6-5 shows the format for this option.

6.1.11 Turn All Reports Off Except Those Specified (Action Code ALLOFF). All reports except those explicitly identified here are to be turned off. The inputs consist of

A B C . . . Y (max of 25)

where A through Y are the report ID numbers (table 6-1) to be turned on. The format is shown in figure 6-6. This option will control the printing of all reports, including histograms, if they contain a specific ID number.

6.1.12 Turn All Reports On Except Those Specified (Action Code ALLON). All reports except those explicitly identified here are to be turned on. The input consists of

A B C . . . Y (max of 25)

A through Y are the report ID numbers (table 6-1) to be turned off. The format is the same as action code ALLOFF (see figure 6-6). This option will control the printing of all reports, including histograms, if they contain a specific ID number.

6.1.13 Continue Data Reduction After an Input Option Error (Action Code ERROR). This code allows data reduction to continue when an error has been detected and reported in an input option request. The default value will abort data reduction and report the error. Only the Action Code card is required.

6.1.14 Debug For a Given Program Number (Action Code DEBUG). This is a debug option which supplies large amounts of output for a given program number. It should be used only in cases of data reduction problems. Card 1 contains the word DEBUG and card 2 contains a program number. A program number of -150 will provide detailed debug on system scheduler activities.

6.1.15 Stop After a Specified Number of Tape Records Processed (Action Code NREC). This option is useful when a tape problem occurs and the entire tape cannot be processed. When this occurs, the program will usually abort with an I/O error and some reports might be lost. If a tape error does occur during data reduction, the operator should type a "U" in

```
Card  1   A
      2   N M
      3   B C D E ...

where
A = The word TIME
N = Report NAME to be time spanned (table 6-1)
M = Number of different times appearing on Card 3
B,C,D,E = Start and stop times used to define the time spans.
          Times must be separated by one or more blanks.
```

Figure 6-5. TIME Action Code Format

6.1.21 <u>Change the Program Number for the First Slave Job (Action Code FSTSLV)</u>. In the GCOS system, certain program numbers are assigned to system jobs. For example $CALC is program number 1, $PALC is program number 2, $SYOT is program number 3, etc. In the WWMCCS system, all programs with a program number less than 14 (decimal) are considered system programs. This option allows the user to alter this program number from its default value of 14. The first card contains the word FSTSLV and the second card contains the new program number. For non-WWMCCS systems, FSTSLV should normally be set to 8.

6.1.22 <u>Request that Certain Jobs be Considered System Jobs (Action Code MASTER)</u>. There are certain jobs executed during the course of a day which have program numbers that would designate these jobs as user jobs (see subsection 6.1.21). However, in actuality they are system jobs and should be considered as system overhead. Examples of such jobs are VIDEO, HEALS, the GMF MONITOR, etc. This option allows the user to define up to ten jobs that should be considered as system jobs. The first card contains the Action Code MASTER. The second card contains the number of jobs to be defined as system jobs. The third card contains the SNUMB of each job to be considered as a system program. Each SNUMB must be followed by at least one blank column. It should be noted that VIDEO, $HEALS, the GMF Monitor and $TRAX will automatically be reported as system jobs and do not need to be requested via this option.

6.1.23 <u>Allocation Status Report Print Control (Action Code PALC)</u>. Due to the excessive amount of output possible from the Allocation Status report, a time control can be set to print only those activities that are in any allocation state greater than the time limit. This time limit defaults to 600 seconds (10 minutes). The first card contains the word PALC and the second card contains the new time limit, in seconds.

6.1.24 <u>Request the Special Job Memory Reports (Action Code SPECL)</u>. If the analyst desires to track the memory demands for a specified number of jobs (not to exceed ten), this input option should be invoked. This option will cause two reports to be produced. One report will indicate every time the requested job(s) was swapped or issued a MME GEMORE/GEMREL for memory, how long it was swapped, or how long the GEMORE was outstanding, and how much memory the job(s) required. In addition, every time the special job issues a MME GEMORE, a line from the Memory Map Report will be generated. This line is generated by default and is not dependent upon whether or not the Memory Map Report is enabled. When the analyst wants to match the Memory Map output to the Special Job output, he must do so based on the time value. For example, if the Special Job Report indicates that FTS issued a MME GEMORE at 16.81057, the user would then examine the Memory Map for a line of output with a time smaller than 16.81057, but where the time on the next line is greater than or equal to 16.81057. For example, the Memory Map might have a line of output with a time indication of 16.81052 where the next line of output was 16.81065. In this case, the line of output at 16.81052 shows what memory looked like at the instant in time that FTS issued the MME GEMORE. If the Special Job

Report indicates that FTS was swapped after issuing the MME GEMORE, the analyst could examine the Memory Map in order to determine why FTS was forced to swap.

A line of the Memory Map is also generated every time the GEMORE for the special job was denied or the special job was forced to swap in order for the GEMORE to be satisfied. This line of the Memory Map would show what memory looked like when the special job was denied the memory request or was swapped from memory. A final line of the Memory Map is produced whenever the special job's memory demand was met. For the swap/denied case and the memory-met case, the Special Job Report and Memory Map are matched by locating identical time values on each report. By generating the Memory Map, the analyst can determine if there are certain jobs that are preventing other jobs from acquiring required memory resources. In this case, the Special Job Report and Memory Map Report can be correlated by matching up the time values from both reports with the identical time values. This is especially useful in an analysis of the Timesharing Subsystem or the File Transfer System.

A second report will also be produced which indicates the average memory size of the job(s) during the course of its execution. This average is taken over increments of time where the time increment used, is the same increment that is used to produce the series of plots. The option consists of three cards where the first card contains the word SPECL, the second contains the number of jobs to be analyzed, and the third card contains the list of SNUMBs separated by at least one blank column.

6.1.25 <u>Process Data on a WW6.4 System (Action Code RN)</u>. If the data reduction program is to be run on a WW6.4 system, the user must use this input option. It consists of the letters RN typed on a data card.

6.1.26 <u>Produce a Memory Map Only Under Certain Memory Demand Conditions (Action Code MAPART)</u>. Due to the enormous amount of output generated by the Memory Map and Out of Core Reports, it is suggested that a site not produce these reports as a standard procedure. However, these reports are very useful in that they do provide a complete picture of memory as well as a total list of all jobs waiting for memory. In order to provide an analyst with the capability of obtaining these reports, without being buried in computer output, this option has been designed. When used, this option states that a line of the Memory Map and Out of Core Reports should be generated only when the number of activities waiting for memory surpasses a certain limit. To invoke this option, a two-card format is required. Card 1 contains the word MAPART and card 2 contains the number of activities that must be waiting memory before a line of output will be generated for the Memory Map and Out of Core Reports.

6.2 Processing

6.2.1 General. The reports of the MUDRP are intended to aid in the following:

- System sizing - both memory sizing and processor utilization.
- Job flow analysis - determining if and where a bottleneck exists and the user memory loading and the daily load distributions.
- System perturbation measures - allows the user to evaluate how a new procedure or new load may alter the utilization of the system as well as determine the total utilization/capacity of the system.
- Large user jobs - aid in identifying specific jobs which are misusing or "hogging" system resources.

Figure 6-7 illustrates how the monitor will pinpoint these various areas. For example, if the monitor indicates a large percentage of processor idleness with high memory demand and low memory availability, a dispatching or I/O bottleneck would be indicated. This would be caused by the I/O not completing its services in a sufficiently timely manner to allow full use of the processors. If processor use was very high and memory demand and availability were high, a memory allocation bottleneck or an overloaded processor would be indicated.

6.2.2 JCL. Figure 6-8 presents the JCL needed to run a total MUM reduction. The following points describe key features of the required JCL.

- 62K required for memory
- 15K sysout requirement would vary depending on amount of data collected. This figure would be significantly higher if the Memory Map or Out of Core Report were produced.
- The DATA I* card is used to indicate the presence of data cards. All data cards must immediately follow this card. At least one data must be present. That card will contain the word END and is used to signify the end of input data. The END card must be present even if no other data cards are desired. The data cards shown in the example are those recommended for most analyses.
- An additional 12K will be required to load the MUM reduction program, but this 12K will be released immediately upon loading.

```
$   IDENT      1820251/30/3044,C702
$   SELECT     B29IDPXO/OBJECT/MUM
$   TAPE       01,X1D,,18897
$   LIMITS     999,62K,-4K,15K
$   DATA       I*
    DATA cards
    ALLON
    14
    5 6 7 12 13 14 15 17 18 29 30 46 39 47
    SPECL
    1
    TS1
    END
```

Figure 6-8. JCL to RUN MUDRP

This figure would indicate what percentage of the total size time produce was used by this program. The size-time product of a job is an attempt to determine the memory effect of a job based not only on its size, but on the length of time that it runs. A 20K jobs that runs for three hours might be more detrimental to a system than a 60K jobs that runs ten minutes.

o $$\frac{\text{Total Size Time Product for This System Program * 100}}{\text{Total Size Time Produce Available to System}}$$

Where Total Size Time Produce Available to System = (The Elapsed Run Time) * (Total Allocatable Memory)

The next two figures are weighted memory sizes for this program. The first figure is the weighted memory size of this program while it is in memory. Therefore, if TSS was in memory during three different time periods for 1/2 hour, 3/4 hour, and 1 hour, and during these periods its memory size was 40K, 100K, 180K respectively, its weighted in memory size would be calculated as follows:

$$\text{Weighted (IN)} = \frac{(40)*(.5)+(100)*(.75)+(180)*(1)}{2.25}$$
$$= \frac{275}{2.25} = 122\text{K}$$

Had the calculation not been weighted by time, the average size of TSS would have been:

$$\frac{(40)+(100)+(180)}{3} = 73\text{K}$$

In the above calculation, the report would be stating that the amount of memory being taken away from the system, by TS, was 122K. However, what if TSS was swapped for 50 percent of the total elapsed time? Then TSS really did not take 122K from the system, but rather only 61K. The second weighted figure takes into accounte the total time the program was actually in memory.

The final figure is the number of times this program was swapped.

In addition to the standard system programs, any jobs requested by the user, to be considered as system jobs, also appear in this report. In figure 6-11 we see 6 user-requested jobs appearing on the report. The user had actually requested nine jobs to be considered as system jobs, but three of those jobs never appeared. In a system using multicopies of TSS only TS1 (prog #5) will appear in this report. Other copies of TSS must be requested by user input option "MASTER". In a WWMCCS system, a program is considered to be a system program if it has a program number less than decimal 14. Commercial users should use the FSTSLV option to change the

THE SYSTEM PROGRAM USAGE OF MEMORY WAS:

PROGRAM		% USED OF MEMORY USE TIME	% OF TIME IN MEMORY	% USED OF SIZE-TIME USED OF MEMORY	% USED OF SIZE-TIME AVAILABLE	WEIGHTED IN/TTM	# SWAPS
(001)	GMF MONITOR	7%	100%	5%	5%	23/ 24	0
(002)	PERIPHERAL ALLOCATOR	5%	71%	2%	2%	17/ 12	988
(003)	GENERAL SYSTEM OUTPUT	6%	93%	6%	5%	26/ 25	58
(004)	REMOTE INPUT	2%	35%	1%	1%	16/ 5	415
(005)	TIME-SHARING EXECUTIVE	7%	100%	31%	28%	129/129	0
(006)	TEST AND DIAGNOSTICS	0%	1%	0%	0%	6/ 0	0
(010)	LOG ON - SECURITY	2%	38%	1%	0%	11/ 4	292
(011)	FILE MANAGEMENT SYSTEM	0%	12%	0%	0%	9/ 1	0
(012)	WW NETWORK CONTROL (NCP)	6%	88%	5%	5%	27/ 24	79
(013)	WW TELNET	4%	65%	2%	2%	16/ 10	150
(014)	WW T2EXEC	3%	47%	3%	3%	28/ 13	233
(015)	WW TLCF	2%	38%	1%	1%	16/ 6	54
(016)	DMS PROCEDURES (DMTEX)	0%	0%	0%	0%	4/ 0	1
SPEC 1	HEALS	1%	21%	0%	0%	7/ 1	150
SPEC 4	NACE	3%	50%	2%	2%	18/ 9	1652
SPEC 6	VIDEO	7%	99%	1%	1%	6/ 5	2
SPEC 7	LGEN	7%	99%	4%	3%	17/ 16	4
SPEC 8	NEWRM	0%	8%	0%	0%	3/ 0	0
SPEC 9	DMSTA	5%	79%	1%	1%	8/ 6	717
					TOTAL	387/290	

Figure 6-11. System Program Load

value of 14 to an 8. In addition, the GMC Monitor, $HEALS, VIDEO and $TRAX are all considered to be system programs.

6.3.3 MUM Reports. The following paragraphs describe the reports output by MUM.

Report numbers 1-50 are all presented in histogram format (see figure 6-12). At the top of the report, the system name, as well as the time and date of data collection, are given. This is followed by the title line of the histogram. Column number 1 indicates the number of occurrences of a given event, with column number 5 describing the event. In figure 6-12, we find that 229 times there were 0 user activities in memory, while 2899 times there were 5 activities in memory. Column number 2 is simply a running total of column number 1. Therefore, the second line in column number 2 (2008) is merely a running total of the first two lines of column number 1 (229 + 1779). The fourth column is the percentage of all activities which will fall into that line of the report showed two user activities in memory. For example, 4063 entries out of a total 23410 entries. This results in a 17.35 percentage figure. This means that 17.35% of all measurements (4063/23410) showed 2 activities being in memory. Column number 3 is simply a running total of column number 4. It presents the percentage of measurements which will fall into a given line, or earlier line. For example, 25.93% of all measurement showed the number of activities in memory to be 2 or less. There is a graphic display of these measurements presented to the right of the fifth column. At the bottom of the report, summary information is provided and is calculated in the standard statistical manner.

In figure 6-13, we see a similar histogram report. As displayed by column 5, we find that each line of the histogram represents a range of values, with an interval size of 200. This interval size can be modified by the user. The lowest value in this histogram is 0 (modifiable by the user) and the size of the histogram defaults to 45 lines (also modifiable by the user). Actually, for this run, the lowest value recorded was 42. Since we can output only 45 lines and each line represents a range of values of 200, the largest value that could be reported would be 9,000 (200 x 45). If a measurement falls outside this maximum value, it is reported as an out-of-range value. In figure 6-13, we find the 21 measurements exceeded 9,000. The average of these 21 measurements was 20188.48. The first line of the summary includes all measurements that were taken. Therefore, 21 out of 79 entries (26 percent) of all measurements were out of range. the average of all measurements taken was 5953.62, while the average of the in-range masurements, (all out of range values are eliminated) was 799.62.

6.3.3.1 Report 1 - Memory Demand Sizes of New Activities in 1K Word Blocks. This report shows the demand size, in 1K-word blocks, of each individual user activity as it was first seen by the memory allocator. The demand sizes are presented to the allocator by the Peripheral Allocator. This is a good measure of the memory demand load of a systems

```
DISTRIBUTION COLLECTED ON SYSTEM NMCC2  AT 12:39:46 ON 80-12-15

                 THE # OF USER ACTIVITIES IN MEMORY

INDIV.  CUMUL.  CUMUL.   INDIV.  NUMBER            PERCENT OF OCCURRENCE                        REPORT
NUMBER  NUMBER   PRC      PRC     WAIT    00   05   10   15   20   25   30   35   40   45   50     15
                                          I....I....I....I....I....I....I....I....I....I....I
  229     229    0.978    0.978    0-   0 IX
 1779    2008    8.578    7.599    1-   1 IXXXXXXXX
 4063    6071   25.933   17.356    2-   2 IXXXXXXXXXXXXXXXXX
 6597   12668   54.114   28.180    3-   3 IXXXXXXXXXXXXXXXXXXXXXXXXXXXX
 7062   19730   84.280   30.167    4-   4 IXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
 2899   22629   96.664   12.384    5-   5 IXXXXXXXXXXXX
  633   23262   99.368    2.704    6-   6 IXXX
  134   23396   99.940    0.572    7-   7 IX
   14   23410  100.000    0.060    8-   8 I

23410 ENTRIES TOTAL    AVERAGE =  3.30145    VARIANCE =   1.622    STANDARD DEVIATION =   1.273
```

Figure 6-12. Standard Histogram Report

```
                    DISTRIBUTION COLLECTED ON SYSTEM NMCC2 AT 18:50:03 ON 81-03-13

                        THE ELAPSED DURATION OF USER ACT IN TENTHS OF A SECOND

 INDIV.   CUMUL.   CUMUL.   INDIV.   TENTHS                 PERCENT OF OCCURRENCE                     REPORT
 NUMBER   NUMBER    PRC      PRC     SECOND   00   05   10   15   20   25   30   35   40   45   50       20
                                              I....I....I....I....I....I....I....I....I....I....I
   22       22    27.848   27.848    42- 199  IXXXXXXXXXXXXXXXXXXXXXXXXXXXX
   14       36    45.570   17.722   200- 399  IXXXXXXXXXXXXXXXXXX
    5       41    51.899    6.329   400- 599  IXXXXXX
    2       43    54.430    2.532   600- 799  IXXX
    1       44    55.696    1.266   800- 999  IX
    0       44    55.696    0.     1000-1199  I
    1       45    56.962    1.266  1200-1399  IX
    2       47    59.494    2.532  1400-1599  IXXX
    1       48    60.759    1.266  1600-1799  IX
    0       48    60.759    0.     1800-1999  I
    1       49    62.025    1.266  2000-2199  IX
    3       52    65.823    3.797  2200-2399  IXXXX
    0       52    65.823    0.     2400-2599  I
    0       52    65.823    0.     2600-2799  I
    1       53    67.089    1.266  2800-2999  IX
    1       54    68.354    1.266  3000-3199  IX
    1       55    69.620    1.266  3200-3399  IX
    0       55    69.620    0.     3400-3599  I
    2       57    72.152    2.532  3600-3799  IXXX
    0       57    72.152    0.     3800-3999  I
    0       57    72.152    0.     4000-4199  I
    1       58    73.418    1.266  4200-4399  IX

  79 ENTRIES TOTAL     AVERAGE = 5953.62024     VARIANCE = 04446990.000     STANDARD DEVIATION = 10219.931
          21(26%) OUT OF RANGE    AVERAGE FOR THESE =      20188.48     IN RANGE AVERAGE =      799.62
```

Figure 6-13. Out-of-Range Histogram

operation and can be used to set System Schedular classes to correctly balance the load cross varying memory size jobs. In this type of report, the distribution shows the percentage of activites which ha a particular memory size. For this report, an entry is made for each new user activity demand at each allocator call. See Report 10 for an explanation of user vs. system activity.

6.3.3.2 Report 2 - The Memory Demand Size of All Demand Types. This report contains the information in Report 1, with the addition of all other individual demand types. These include activities that are swapped or involved in a memory compaction procedure. This report should be similar to Report 1, unless a great amount of GEMORE, GERLEC, or swap operations are performed by the users load. This would alter the memory size demands from that seen by the allocator at the initial request. For this report, an entry is made for each activity with an outstanding demand for each allocator call. Activities with an urgency of 0 are not counted.

6.3.3.3 Report 3 - The Total Memory Demand Outstanding. This report shows the sum of demand for all activities in the system including outstanding GEMOREs. It is a distribution of memory demand that is not satisfied, across the measurement session. It should be remembered that all data is collected at the Core Allocator and does not represent the full system load. Portions of the load may be held in the System Schedular and the Peripheral Allocator. Activities with an urgency of 0 are not counted. For this report, an entry is made at each allocator call. For most analyses, this report will not be used since report 7 provides a more statistically accurate representation of this data.

6.3.3.4 Report 4 - The Demand That Was Outstanding When a Processor Went Idle. This report is the same as Report 3, except that an entry is made only if a processor has gone idle since the last allocator call. If a large demand should be outstanding during processor idleness, a system bottleneck may be present. In this case, memory is probably fully utilized (i.e., demand cannot be satisfied), but the activities that are occupying memory are not using the processor, (i.e., a processor has gone idle). This is a good sign of an I/O backlog. IDLEM data is used to produce this report. If the Idle Monitor was not active, this report will not be produced.

6.3.3.5 Report 5 - The Total Amount of Available Memory. The total amount of available memory is a key indicator of the system memory utilization. If this amount is continually low, the memory is being fully utilized and possibly in need of expansion. A continually high amount may indicate another system bottleneck or an excess of memory. This report, when used in conjunction with Reports 3, 4, and 6 should give a good first-level indication of system memory utilization. It should be noted that the availability shown here exists in all quadrants. The availability is the sum of any and all "holes" in the system and does not mean that this memory is contiguously available.

The average value reported in this report minus the average value reported in report 3 will give a good feel for memory surplus or shortfall. A positive result will indicate a surplus while a negative result will indicate a shortfall. The MUM heading report also gives a surplus/shortfall indicator. Any activity with an urgency of 0 that is currently in memory will have its memory size included in this availability figure. The reason for this is that if memory becomes a constraint, these activities can be swapped and their memory will become available for use.

For this report, an entry is made for each allocator call. For most analyses, this report will not be used since report 8 provides a more statistically accurate representation of this data.

6.3.3.6 Report 6 - The Memory Available When a Processor Went Idle. The previous report is repeated with the additional restraint that a processor has gone idle since the last allocator call. This aids in identifying either a bottleneck or a lightly loaded system.

For this report, an entry is made at each allocator call that had a processor go idle since the last allocator call. IDLEM data is used to produce this report. This report will not be produced if IDLEM was not active or the IDLEM Reports have been disabled via user input command.

6.3.3.7 Report 7 - The Time-Corrected Total Demand Outstanding. See report 16 for an explanation of time correction. The time-corrected total demand is the sum of all requests for memory known to the allocator as indicated in report 3. Activities with urgency 0 are not counted.

6.3.3.8 Report 8 - The Time-Corrected Memory Available. See report 16 for an explanation of time correction. This report reflects the time-corrected amount of total memory available as indicated in report 5.

6.3.3.9 Report 9 - The Number of Activities Waiting for Memory in Allocator Queue. This report identifies the depth of the allocator demand queue and includes all activites that are waiting for memory allocation. Activities with a 0 urgency are not considered as waiting for memory. This report aids in determining if too many or too few activities are getting to the Core Allocator from the Peripheral Allocator. For this report, an entry is made at each allocator call. For most analyses, this report will not be used since report 11 provides a more statistically accurate representation of this data.

6.3.3.10 Report 10 - The Number of User Activities Waiting Memory in Allocator Queue. This report is the same as report 9 except that it only counts those activites of a slave job as identified by their program number (program number 14 or greater). In order to change this program number test, the user should see Input Action FSTSLV. In addition, the user may specify up to ten additional programs that he wants considered as system programs, even though their program number exceeds 14. The user

should see Input Action MASTER in order to select this option. This report indicates the "user" work waiting allocation. For this report, an entry is made on each allocator call. For most analyses, this report will not be used since report 12 provides a more statistically accurate representation of this data.

6.3.3.11 Report 11 - The Time-Corrected Number of Activities Waiting Memory. See report 16 for an explanation of time correction. This report indicates the time-corrected number of activites waiting memory as in report 9.

6.3.3.12 Report 12 - The Time-Corrected Number of User Activities Waiting Memory. See report 16 for an explanation of time correction. This report indicates the time-corrected number of user jobs waiting memory in the allocators queue as in report 10. See report 10 for additional user options.

6.3.3.13 Report 13 - The Number of Activities Waiting Memory When a Processor Went Idle. Report 9 is the basis for this report, with the additional criteria that a processor must have gone idle since the last allocator call. An entry is made for each allocation where a processor has gone idle since the last call. IDLEM data is used to produce this report. This report will not be produced if IDLEM is not active or the IDLEM reports were disabled via user input commands.

6.3.3.14 Report 14 - The Number of Activites Residing in Memory. This report represents the number of activities allocated memory. It indicates the multiprogramming depth the system is obtaining. It is probably an upper level since an activity is allocated memory prior to and past actual usage. Any activity in memory, with a 0 urgency, is not considered as residing in memory. For this report, an entry is made for each allocator call. For most analyses, this report will not be used since report 16 provides a more statistically accurate representation of this data.

6.3.3.15 Report 15 - The Number of User Activities in Memory. The activites shown in this report are those that are in memory and have a program number greater than or equal to 14. These are user programs. For this report, an entry is made at each alloator call. See report 10 for additional user options in defining system jobs and user jobs. For most analyses, this report will not be used since report 17 provides a more statistically accurate representation of this data.

6.3.3.16 Report 16 - The Time-Corrected Number of Activities in Memory. This report presents the same information as in report 14. The number of entries at each allocator call is determined by the time since the last allocator call. The result is a simulation of a uniform sample rate of allocator calls. Therefore, the noncorrected reports display the distributions as seen by the allocator itself. The time-corrected reports present the time weighted distributions. As an example assume that three

measurements are taken. It is found that 6 activities are in memory for 2 minutes, 20 activites for 5 minutes, and 8 activities for 1 minute. The average number of activites in memory is (6+20+8)/3=11. If we correct for time however, we get ((6)*(2)+(8)+(20)*(5))/8=100/8=12.5 activites in memory. The division of 8 was the total time (5+2+1) spent collecting data. All of the time-corrected reports are of the same nature.

6.3.3.17 Report 17 - The Time-Corrected Number of User Activities in Memory. This report indicates the time-corrected number of user jobs with allocated memory as in report 15. See report 10 for additional user options in defining system jobs and user jobs. See report 16 for a definition of Time Correction.

6.3.3.18 Report 18 - The Number of Activities in Memory When a Processor Went Idle. This report indicates the total number of activities with allocated memory when a processor went idle. This report can show an I/O bottleneck if the multiprogramming depth is high but there is no work for a procesor to perform. For this report, an entry is made on each allocator call for which a processor went idle since the last call. IDLEM data is used to produce this report. This report will be not be produced if IDLEM is not active or if IDLEM reports have been disabled by user input options.

6.3.3.19 Report 19 - The Ratio of User Activity Duration Versus Its Memory Use Time. This report indicates the ratio of total elapsed time (report 20) over the total allocated memory time (report 21). This shows how activity run time is stretched due to memory resource contention. If there was no memory contention then the total elapsed time of an activity would be very close to total memory time of an activity. As memory contention increases (i.e., swapping occurs), the total elapsed time will begin to increase in comparison to total memory time. It must be realized that total memory time can increase if there is CPU or I/O contention. As a job sits in memory, waiting on the CPU or I/O subsystems, its memory time will increase.

For this report, an entry is made for each user activity that terminates. See report 10 for an explanation of user vs. system activities.

6.3.3.20 Report 20 - The Elapsed Duration of User Activity in 10ths of a Second. This report presents the clock time that the allocator knew of a user activity's existence, measured from its first memory demand to its termination. This includes all time spent in a GEWAKE, in memory, and swapped.

For this report, an entry is made for each user activity that terminates. See report 10 for an explanation of user vs. system activities.

6.3.3.21 Report 21 - The Total Elapsed Time a User Activity Was in Memory. This report shows the duration of elapsed clock time each user

activity had memory allocated to it. It helps describe the system workload requirements.

For this report, an entry is made for each user activity that terminates. See report 10 for an explanation of user vs. system activities.

6.3.3.22 Report 22 - The GEMORE Service or Denial Time - 1/10 Second, Elapsed. The time from a GEMORE request until the activity is allocated the extra memory, swapped to achieve the aditional memory, or denied the memory is displayed in this report.

For this report, an entry is made for each activity whose GEMORE request is no longer present. This report is not used in most analyses.

6.3.3.23 Report 23 - The Request Size of GEMOREs. All GEMORE requests are shown in this report with the displayed size in 1K blocks.

For this report, an entry is made for each GEMORE request. This report is not used in most analyses.

6.3.3.24 Report 24 - Delay Time in the System Scheduler. The amount of time a job spends in one of the scheduler queues is displayed in this report. The Allocation Status Report can be used to display particular jobs that are delayed for excessive amounts of time.

6.3.3.25 Report 25 - Delay Time Until Core Allocation. This report displays the total amount of time activities spent in the various allocation phases prior to core allocation. The Allocation Status Report can be used to display particular activities that are delayed for excessive periods of time.

6.3.3.26 Reports 26 through 31 - The Elapsed Time of a Busy State of the Processors. These reports present the elapsed clock time between the idle states of each individual processor. The reports supply an indication of how each processor is utilized versus the others in the system.

For these reports, an entry is made at each idle state of a processor. IDLEM data is used to produce these reports. These reports will not be produced if the IDLEM was not active or if the IDLEM reports have been disabled by user input option.

6.3.3.27 Report 32 - The Elapsed Time of a Busy State of Processors. The elapsed clock time between idle states of all processors is presented in this report.

For this report, an entry is made for each processor idle state. IDLEM data is used to produce this report.

6.3.3.28 Report 33 - Elapsed Time Between Allocator Calls in 1/100 of a Second. This report shows the elapsed clock time between calls to the allocator and shows if the allocator is receiving sufficient service.

For this report, an entry is made for each user activity that terminates. This report will not be used for most analyses.

6.3.3.29 Report 34 - The I/O Time Charged per User Activity in Seconds. This report indicates the I/O time charged to each user activity.

For this report, an entry is made for each user activity that terminates.

6.3.3.30 Report 35 - The CP Time Charged per User Activity in Seconds. This report presents the CP time charged to each user activity. For this report, an entry is made for each user activity that terminates.

Reports 34 and 35 report the total CPU and I/O times used by a user activity while the monitor was active. These histograms are not generated for programs with program numbers less than 14 (i.e., system programs). See report 10 for additional user options in defining system activities and user activities.

6.3.3.31 Report 36 - The Number of Times a User Activity was Swapped. This report shows the swap count per user activity. The total number of swaps a user activity incurs is the user argument, as counted by the monitor. See report 10 for additional user options in defining system activities and user activities.

For this report, an entry is made for each user activity that terminates.

6.3.3.32 Report 37 - The Total Elapsed Time a User Activity was Swapped. This report indicates the total time a user activity was inactive due to a swap. After each swap is completed, an accumulator is updated, and if an activity is terminated, an entry is made to this report. See report 10 for additional user options in defining system activities and user activities.

6.3.3.33 Report 38 - The Number of Times a System Activity was Swapped. This report is the same as report 34 except for system activities. See report 10 for additional user options in defining system activities and user activities.

6.3.3.34 Report 39 - The Total Elapsed Time a System Activity was Swapped. This report is the same as report 35 except for system activites. See report 10 for additional user options in defining system activities and user activities.

6.3.3.35 Report 40 - Number of Extra Activities That Might Fit in Memory Using Compaction. This report shows how memory might have been used more optimally. It takes the total amount of available memory (displayed in

report 5) and attempts to fit in those activities waiting memory. If an activity fits, the memory available is decreased, and the next activity is tried. If an activity does not fully fit, the next activity is tried. This continues until all available memory is used or until all the activities waiting have been tried. The search starts at th first waiting program an progresses serially down the program numbers of those waiting. This search ignores the actual size of "holes" or quadrant-crossing and is not necessarily obtainable or optimal. For this report, an entry is made at each allocator call.

6.3.3.36 Report 41 - Number of Extra Activities that Might Fit Memory Without Compaction. This report is the same type as report 40. In this case, activities are fit into existing holes and are ordered by urgency. The search progresses down the activities serially, beginning at the highest urgency activity. This histogram presents a good picture of how well the core allocator is performing its function.

For this report, an entry is made at each allocator call.

6.3.3.37 Report 42 - The Percent of Size-Time Product Used by a User Activity. This report shows the percentage of users each activities' size-time product over its run-time duration. An entry is made for each user activity that terminates. See report 10 for an explanation of user and system activities.

6.3.3.38 Report 43 through 49 - The Length of Idle State in the Processors. The elapsed clock time of an idle state is given in these reports for each individual processor and also as an average for all processors. They supply an indication of how each processor was utilized versus the others in the system. They also provide information on how busy the processors are. These reports should be used in conjunction with reports 26 through 32.

DEVICE SEEK MOVEMENT SUMMARY FOR SYSTEM NMCC2 ON 82-06-10

DEVICE LOCATION	DEVICE NAME	CONNECTS	AVERAGE SEEK LENGTH
SEEK MOVEMENT OF IOM-0,PUB-08,DEVICE-01-- MS0450	ST1	33559	56.20373
SEEK MOVEMENT OF IOM-0,PUB-08,DEVICE-02-- MS0450	DC2	8505	76.86820
SEEK MOVEMENT OF IOM-0,PUB-08,DEVICE-03-- MS0450	DC3	4630	68.88812
SEEK MOVEMENT OF IOM-0,PUB-08,DEVICE-04-- MS0450	DC4	5979	51.75682
SEEK MOVEMENT OF IOM-0,PUB-08,DEVICE-05-- MS0450	RF5	4213	37.23190
SEEK MOVEMENT OF IOM-0,PUB-08,DEVICE-07-- MS0450	RF7	278	0.00719

Figure 7-24.3. Device Seek Movement Summary Report

processors and provides the processor number involved as well as the time of day the assignment/releasing of processors occurred.

- FOLLOWING PRINTS ARE THE INPUT OPTIONS . . .
  An echo print of all nonstandard input options selected will be produced. If any input option is described incorrectly, an error message will be generated indicating the type of error and the card number in error. The user should correct the error and resubmit the job for processing.

- FOR INFORMATION ONLY
  This message will then be followed by several lines of output describing special record types that have been processed, or special processing events that have been executed by the MSMDRP. In most cases, the message can be ignored. Those messages which are important, and reveal an error in processing logic, will be described below. All other messages will not be described.

- JULIAN DATES ARE BAD - RUN TERMINATED
  Every GMF data record is preceded by the current Julian date. The MSMDRP has found a Julian date that does not agree with the Julian date found on previous records. This can occur when an old GMF data tape is reused without degaussing. Old data is on the tape, and if the new data failed to write an end of file mark on the tape because of a system crash or malfunction, the MSMDRP, after reaching the end of the new data, would attempt to process the old data without realizing that it was old. The check on the Julian date prevents this from happening. The MSMDRP will terminate cleanly and all reports will be produced.

- HAVE INCREASING OR BAD SEQUENCE NUMBERS . . .
  A problem has occurred in reading the data tape. If the run is reprocessed, the error may disappear. If it reoccurs, then the tape was generated with an error. In most cases, the MSMDRP will recover and processing will not be significantly affected. If it occurs often, contact CCTC/C751.

- PROCESSING TERMINATED BY NXTRECRD . . .
  MSMDRP has requested the operator to mount a new tape and the operator responded that he did mount the new tape. However, MSMDRP is unable to match the inital record on the new tape with the last record on the old tape. User should check the data collection procedure to insure that correct tapes were mounted during the data collection phase. MSMDRP will terminate cleanly and all reports will be produced.

- INCURRED A BAD SEEK ADDRESS . . .
  In the logic processing of subroutine SYSTMFIL, an unexpected condition has occurred. MSMDRP will continue processing correctly, but if this occurs frequently, CCTC/C751 would like to be notified. Call 202-695-0856.

7.6 Default Option Alteration

Most users rely upon the standard MSM Report formats and their default values as these suit a wide range of needs. A capability to change the reports is built into MSMDRP. The general form for all option requests are as follows: The first card contains an action code describing the action to be taken. Subsequent cards modify report parameters for some of the action codes. All input cards are free format with the only requirement being that at least one blank space separates multiple input parameters. The very last input card must have the word "END" entered in it. This card must be present whether or not any other input options are selected.

There is no specific order required of the options, and multiple entries of each are permissible. If several inputs refer to the same report, the last one encountered will have precedence. If a report is turned off by default and is modified, it will be turned on through the request for modification. The chart below shows the available actions: the mnemonic code for the user to identify the action; the function; and the default.

Mnemonic	Function	Default (indicated in parentheses)
AREA	Request file code references made to a specific area of a specific device	(not provided)
DEBUG	Debug	(no debug)
ERROR	Do not stop on Input Error	(stop)
FILDEF	Define system files by name	(no names used)
END	This card must be present	
MODULE	Produce the SSA Module Usage Report by Job	(no report produced)
NCONN	Process a limited number of connects	(total tape processed)
NREC	Process a limited number of tape records	(total tape processed)
OFF	Turn reports off	(all reports ON except reports 12,16, 18,20 - see table 7-1)
ON	Turn reports on	(all reports on except reports 12,16, 18,20 - see table 7-1)
PROJ	Produce the Connect Summary Report by Userid/SNUMB	(no report produced)

RN	This option must be selected when .MSMDRP is used to process WW6.4 data or when MSMDRP is executed on a WW6.4 system	(process WW7.2/4JS data on a WW7.2/4JS system)
TIME	Set a timespan for measurement	(no time criterion)
TIMEQ	Change time quantum for Chronological Device Utilization Report	(report is off - default value is 60 seconds)
USERID	Suppress userids from reports	(userids printed)
RATECH	Change time quantum for Connects Per 10 Minute Report	(report is off - default value is 10 minutes)
CAT	Turn on the Cat/File String Report	(report off)
RATE	Request the Connect Per 10 Minute Report for specific user jobs.	
LIMITS	Limit the amount of processing performed and reports produced.	

7.6.1 Monitor a Specific Device Area (Action Code AREA). This option allows a user to specify specific areas of a device for which all jobs referencing this area are to be highlighted. The format of the display is that of a File Code Summary and contains those jobs and file codes that reference the area of interest.

The device to be investigated is identified via the PUB and IOM number. The specific areas of interest are identified as beginning at the starting address defined in llinks. The length of the area is also in llinks, with a zero meaning the end of the device. A total of ten possible areas are allowed. The format for this card is shown in figure 7-25.

See subsections 7.5.12 and 7.5.16 for complete details on the report format generated with this user option. This report is off by default and will be activated by the processing of this action code.

7.6.2 System Debug (Action Code DEBUG). This is a restricted option for GMF system developers. DEBUG should only be used with guidance received by CCTC/C751.

7.6.3 Continue Data Reduction After an Input Option Error (Action Code ERROR). This option allows data reduction to continue when an error has been detected and reported in an input option request. The default value reports the error and aborts the data reduction procedures. The format for this option is the word ERROR on the data card.

7.6.4 Specify System File Names (Action Code FILDEF). This option allows the user to specify the name of each system file displayed in the

Card 1 = A
Card 2 = N
Card 3 = B C D E F

Where

A = The word AREA
N = The number of areas to be specified. A maximum of ten areas are permitted. A Card #3 must be present for each area requestred.
B = IOM number
C = Pub number
D = Device number
E = Starting address in llinks
F = Length of area in llinks

The following definitions apply to this option.

Device Type	Numbers Cylinders	Number Sectors/ Cylinder	Number Sectors/ Block (LLINK)
180	200	360	5
181	200	360	5
190	407	589	5
191	407	760	5
450	811	760	5

Figure 7-25. Specific Device Area Report Card Input

System File Use Summary Report discussed in subsection 7.5.9. This option is specified with a set of data cards. The first data card contains the word FILDEF. The second data card contains the number of system files to be described on the following cards. The following cards each contain a single pair of data points separated by at least one blank. The first data point is the system file number and the second data point is the desired system file name.

The standard output of the System File Use Summary Report is to label each system file as System File 1, System File 2, etc., corresponding to GCOS-HI-USE, GCOS-LO-USE, etc. In order to know the correct order of the file names, the user should check the $ FILES section of the startup deck. The order of the files in the $ FILES section of the startup deck is the order they are referenced in the report.

7.6.5 End Card (Action Code END). This card must be present at all times and must be the last data card supplied. It consists of the word END entered on the card.

7.6.6 Produce the SSA Module Usage Report by Job (Action Code MODULE). This option allows the user to produce the SSA Module Usage Report. This report will list every SSA module used by every job that was run during the monitoring session. See subsection 7.5.11 for details concerning this report. This report is off by default and cannot be turned on by using the ON option. This report can be activated only by entering MODULE on the data card.

7.6.7 Record Limitation by Connects (Action Code NCONN). This option allows a user to process only a specific number of connects. This option is especially useful if the tape contains an error on it and cannot be completely processed. Using this option, the user can process the tape or tapes up to the point where the tape error exists. This option requires two data cards. The first data card contains the word NCONN with the second card containing the number of connects to be processed.

7.6.8 Record Limitation by Records (Action Code NREC). This option allows a user to process only a specific number of tape records. This option is especially useful if the tape contains an error on it and cannot be completely processed. Using this option, the user can process the tape or tapes up to the point where the tape error exists. This option requires two data cards. The first data card contains the word NREC with the second card containing the number of tape records to be processed.

7.6.9 Turn a Report Off (Action Code OFF). This option allows a user to turn a report off that is on by default. In MSMDRP, all reports are on except report numbers 12, 16, 18 and 20 (see table 7-1). Only those reports in table 7-1 that have a name in parentheses () can be turned off with this option. Two data cards are required to use this option. The first card contains the word OFF and the second card contains the name of the report as displayed in the parentheses () in table 7-1.

7.6.10 Turn a Report On (Action Code ON). This option allows a user to turn a report on that is off by default. In MSMDRP, all reports are on except report numbers 12, 16, 18 and 20 (see table 7-1). Only those reports in table 7-1 that have a name in parentheses () can be turned on with this option (#9,10,12,15,16,17,18,20). Two data cards are required to use this option. The first card contains the word ON and the second card contains the name of the report as displayed in the parentheses () in table 7-1.

7.6.11 Produce Connect Summary Report by Userid/SNUMB (Action Code PROJ). This option allows the user to specify up to a total of 40 USERIDs and SNUMBs for which he wants the Connect Summary Report by Userid/SNUMB produced. The number of SNUMBs requested cannot exceed 10. In addition, the user can request the entire File Code Summary Report, or the user may want to see the File Code Summary Report only for a prespecified set of jobs or USERIDs, or the user may not want the File Code Summary Report at all. For example, the user can request 35 different USERIDs and 5 SNUMBs or 40 different USERIDs and 0 SNUMBs or 30 different USERIDs and 10 SNUMBs or 3 different USERIDs and 6 SNUMBs, etc. The format for this option is shown in figure 7-26. If values of zero are desired, they must be punched on the card. A blank is not equivalent to a zero. The Connect Summary Report will indicate for each requested USERID or SNUMB the number of connects made by the job or USERID. If a requested SNUMB also has a requested USERID, the number of connects issued by that job will be reported twice in the summary report. Refer to subsection 7.5.14 for a description of the report to be produced with this option. If the user desires to see a File Code Summary Report, it will be turned on via this option. The user does not need to use the "ON" input option.

7.6.12 Reduce WW6.4 Data or Process MSMDRP on a WW6.4 System (Action Code RN). This option requires two cards. The first card has the letters RN and the second card one of the following numbers:

1 - WW6.4/2H system processing WW6.4/2H data
2 - WW6.4/2H system processing WW7.2/4JS data
3 - WW7.2/4JS system processing WW6.4/2H data

The default is a WW7.2/4JS system with WW7.2/4JS data.

7.6.13 Set a Timespan of Measurement (Action Code TIME). The timespan of data collection can cover many hours of which only a few may be of interest. This option allows a user to specify the timespan (or spans) to displayed in all reports. For example, the user may specify that he wants to collect data from 0500 to 2200 and wants to display data only from 0900 to 1700 in all reports.

If the entire reduction will have a set timespan, the name "TOTAL" is used. Histogram reports cannot be individually timespanned. All timespans of "other" reports will be bounded by the overall report timespan, if one will be used. Up to five timespans for each report type may be specified,

Card 1 = A
Card 2 = B C D
Card 3+ = E
Card 4+ = F

Where

A = The word PROJ
B = 0 if Connect Summary Report is desired, but no File Code Summary Report is desired. The 0 must be punched on the card. A blank is not equivalent to a 0.
B = 1 if Connect Summary Report is desired and a complete File Code Summary Report is wanted
B = 2 if Connect Summary Report is desired and only a partial File Code Summary Report is wanted
C = Number of Userids (30 MAX)
D = Number of SNUMBS (10 MAX)
E = A total of C Userids with one Userid per card
F = A total of D SNUMBS separated by at least one blank space. All SNUMBS should fit on one card.

Figure 7-26. Limited File Code Summary Input Card Format

and they must be serially ordered. Names for other reports may be found in table 7-1. Only those reports in table 7-1 that have the NAME=specification can be time controlled. If a report does not have this specification, it cannot be time controlled. If the entire reduction is to be time spanned, the name that should be entered on the data card is TOTAL. The format for this action is given in figure 7-27. All times are expressed as four character fields with no intervening blanks. Time is based on a 24-hour clock. If the user wants to request the time 4:07 he must input a "0407".

If a start time but no stop time is desired, no characters should be entered after the minutes of the start time. If stop time is requested, there must be a start time corresponding to it.

The File Code Summary Report, the Activity Summary Report and the Device Area File Code Reference Report are all considered as a single unit under this option. Whenever a time frame is reached for any of these reports, all reports will be spanned. As an example, suppose that the user requested the following:

- o File Code Summary Report for 1500-1600 and 1700-1800.
- o Device Area File Code Reference Report for 1530-1730.
- o Activity Summary Report for 1400-1530

For the above requests, the following report spans would be produced:

- o At 1530, the File Code Summary Report would be produced for the period 1500-1530 and the Activity Summary Report would be produced for the period 1400-1530.
- o At 1600, the File Code Summary Report would be produced for the period 1530-1600, and the Device Area File Code Reference Report would be produced for the period 1530-1600.
- o At 1730, the File Code Summary Report would be produced for the period 1700-1730 and the Device Area File Code Reference Report would be produced for the period 1600-1730.
- o Finally, at 1800, the File Code Summary report would be produced for the period 1730-1800.

7.6.14 Change the Time Quantum Value For the Chronological Device Utilization Report (Action Code TIMEQ). The user can change the time quantum value used to produce the Chronological Device Utilization Report by inputting the quantum in seconds. The default value is 60 seconds. Two cards are required. The first card contains the word TIMEQ and the second card contains the new quantum in seconds.

7.6.15 Suppress the USERIDs (Action Code USERID). The user can suppress the printing of USERIDs on the File Code Summary Report by entering the word USERID on a data card.

Card 1 A
Card 2 B
Card 3 C D E F . . .

Where

A = The word TIME
B = Number of different times appearing on Card 3
C,D,E = Time used to define a timespan. Individual times must be separated from each other by at least one blank column. All times are considered to be on a 24-hour clock and must be expressed as a 4-digit field.

Figure 7-27. Input Option TIME Card Format

7.6.16 Change the Time Quantum Value for the Connect Per 10 Minute Report (Action Code RATECH). The user can change the time quantum value used to produce the Connect Per 10 Minute Report by inputting the quantum in seconds. Two cards are required. The first card contains the word RATECH and the second card contains the new quantum in minutes. The default value is 10 minutes.

7.6.17 Turn on the Cat/File String Report (Action Code CAT). This option, consisting of the word CAT on the data card, will turn on the Cat/File String Report (see subsection 7.5.13).

7.6.18 Request the Connect Per 10 Minute Report for Specific User Job (Action Code RATE). This option will allow the user to obtain the Connect Report for specific jobs as well as for the TimeSharing Subsystem and the Total System (see subsection 7.5.24). Card number 1 contains the word RATE, card number 2 the number of jobs desired (a maximum of 8 are permitted), and card number 3 the SNUMBs of the jobs desired. In addition to the requested jobs, the TimeSharing Subsystem as well as the Total System will also be reported. If multiple copies of TSS are in use, all activity will be reported under the single job name of TS1. If the user wants to obtain this report for only TimeSharing and the Total System, then he simply needs to use the "ON" input option using the name "RATE" for the required report ID.

7.6.19 Limit the Processing and Output (Action Code LIMITS). This option will allow the user to control the amount of output produced and the amount of record processing performed. Card number 1 contains the word LIMITS and card number 2 contains either the word ONLYSP or the word NOHIST or the word SUMARY. If the word ONLYSP is used then the Mass Store Monitor program will process only those data records that are generated by the SNUMBs requested under the RATE input option (see subsection 7.6.18). All other data will be ignored. The user must take care when examining the histograms and reports that are produced. The user must remember that only a limited amount of data has been processed. If the word NOHIST is used then no seek or space utilization histograms will be produced. This option can be used in conjunction with the ONLYSP option (must have two LIMITS input cards) or can be used by itself. In the latter case, all data will be analyzed, but no histograms will be produced. Finally, the user can request a summary of the seek movement activity. He can obtain this summary whether or not he selects to produce the set of histograms. To obtain the summary report, he must type SUMARY on a card immediately following the LIMITS card. A summary listing will not be produced for the space histograms, as this summary information is meaningless for this set of histograms.

7.7 JCL

The data reduction procedures consist of a single FORTRAN program having a main level and multiple subroutines.

A description of the more important JCL cards is presented below (see figure 7-28).

The $:LIMITS card should be studied to meet user needs. The run time (99) and output limit (30K) may both need to be altered as required by the duration of the monitoring run. The MSMDRP requires 55K of memory in order to execute plus an additional 2K for SSA space. During the initial loading process, MSMDRP will actually require 68K of memory, but 11K will be released immediately upon loading.

The statement:

```
$    DATA        I*
```

is used to identify the data cards that follow as described in subsection 7.6. At least one data card is required, that being an "END" request.

7.8 Multireel Processing

If more than a single reel of data has been collected, a series of messages will be outputted to the console informing the operator that a new data reel is required. The following are the messages produced.

a. DISMOUNT REEL #XXXXX THEN MOUNT REEL NUMBER YYYYY ON ZZZZZ

   In this case, XXXXX is the old reel number, YYYYY is the new reel number, and ZZZZZ is the tape drive ID.

   If the operator fails to mount the new tape, the above message will be repeated three times, after which the program will terminate, and all reports produced.

b. IS TAPE XXXXX MOUNTED ON DRIVE ID YYYYY (Y/N)

   In this case, XXXXX is the tape number being requested for mounting and YYYYY is the tape drive ID.

   This message occurs when the data reduction program finds the wrong tape has been mounted (by comparing internally generated tape labels). If the operator answers N, the message in (c) below is produced. If the operator answers Y, the data reduction program will terminate and all reports will be produced. In this case, the data reduction program is unable to process the tape. Even though the operator is mounting the correct tape, the internal label on the new tape does not match that being requested by the old tape. The user should check the data collection session to insure that the operator did not respond with an incorrect tape number during multireel change.

```
Col 1   8        16
$       IDENT    1820251/30/3044
$       SELECT   B29IDPXO/OBJECT/MSM
$       TAPE     01,X1D,,12345
$       LIMITS   99,55K,-4K,30K
$       DATA     I*
        Data cards - at least an "END" card must be present
$       ENDJOB
```

Figure 7-28. DRP MSM JCL

After entering the Y or N, the operator will need to hit the EOM key twice in order for the response to be transmitted.

c. WRONG REEL JUST MOUNTED, DISMOUNT AND MOUNT REEL XXXXX ON ZZZZZ

In this case, XXXXX is the new reel number, and ZZZZZ is the tape drive ID.

d. CAN TAPE XXXXX BE MOUNTED ON DRIVE YYYYY (Y/N)

In this case, XXXXX is the new desired reel and YYYYY is the tape drive ID.

If the operator fails to answer this message it will be repeated until he responds with a "Y" for YES or "N" for NO. If he types in "Y", then message (a) will be repeated. If the types in "N", then the program will be terminated and all reports will be produced.

7.9 Tape Error Aborts

During the course of processing it is possible that the operator will be required to abort the data reduction program due to an irrecoverable tape error. If such a condition occurs, the operator should abort the job with a "U" abort. This will allow the data reduction program to enter its wrap-up code processing and produce all reports generated prior to the tape error.

OPCODE	DESCRIPTION
11	Output not available
16	Reject request (temporary)
17	Reject request (permanent)
20	Terminal rejected
110	Backspace output

These opcodes indicate a delay in data transmission or a communications problem. If these opcodes show up consistently, and in significant numbers, a detailed analysis should be conducted.

9.5.10 Response Time Report. This report is produced whenever the user sets the interval time using the input option RATE (subsection 9.6.11) or SNUMB (subsection 9.6.12). The report shows, for each interval, the time of day, the response time in general (i.e., averaged over all DAC subsystems), the response time for the requested subsystems, and the number of opcode rejects, RSVPs and RJMs. (See figure 9-16). The column headings are as follows:

TOD - Time Of Day

RESP - average response time over the time period

I/R - average response time of those responses considered acceptable (see sections 9.5.6 and 9.6.6)

#I - number of responses in the acceptable range

#O - number of responses in the unacceptable range. This number is important in validating the figure in the RESP columns. One extremely bad response can cause a skewed average response.

USER - average number of users on this subsystem during the time frame

OPREJT - number of Opcode Reject Temporary commands received during the period

OPREJP - number of Opcode Reject Permanent commands received during the period

RJM - number of Reject Message commands received during the period

RSVP - number of Resend requests received during the period

NOTE: If TSS is one of the SNUMBs requested, all TS SNUMBs (TS1-TS4) will be represented under TSS.

9.5.11 Error Messages. The CAMDRP can produce multiple error messages relating to the data type. Most of these messages are actually warning messages, which the CAMDRP will try to recover from and will continue to process.

The most prevalant error message is the warning message "TERMINAL ID NOT FOUND." This message usually occurs when a terminal has been logged onto the system prior to the CAM starting to collect its data. When the CAMDRP tries to find a particular user who is receiving or transmitting data in its tables, that user will not be found since the CAMDRP did not find any log-on record for him. The user is logged onto the system and the CAMDRP continues processing.

The main reason for the CAMDRP to abnormally stop processing is the error message "NO MORE ROOM IN TERMID ARRAY." This means that an internal array has been filled. This is usually the terminal ID array. The parameter MAX must be increased to enlarge the required arrays. The current size of MAX is 50. This can be exceeded if there are a large number of users on the system when the CAM is started. To increase this value, the user should log onto TS1. Enter EDIT O B29IDPXO/SOURCE/CAM. Then enter:

```
CASE
VERI
RS:/MAXb=b50/;*:/MAXb=bxx/
```

where xx is the new maximum number of terminals to be allowed on the system. The CAM data reduction program will be increased by 175 words for each terminal above 50 allowed. All other messages produced will be self-explanatory. If they do not indicate a severe error, the words "For Information only" will appear with the message.

9.5.12 H6000-DN355 Reject Report. This report displays all the terminal IDs that have had some type of error opcode from or to the DN355. These opcodes are RJM, RSVP, ECHO, OPCRJT, OPCRJP (see figure 9-16.1).

9.5.13 Abort Report. This report indicates each time the DN355 aborts and is reinitialized. It also presents each time line IDs 01,02 disconnect. These disconnects indicate a WIN problem since WIN has lost its network connection (see figure 9-16.2).

9.5.14 TS1 Initial Parameter Report. This report indicates the initial preset values for TS1. These values are SIZE parameters, LIMIT parameters, SWAP FILES parameters and a list of all ALLOCATED DEVICES per file code. This report is produced once, so if any parameters are changed during the run (such as TS1 max size), the change is not reported. See figure 9-16.3 for a sample of this report.

9.5.15 Mailbox Busy Report. This report is printed out each time a Response Time Report line is printed. This report indicates the running total of special interrupts that have occurred during the time frame and the average number of unanswered interrupts, requests waiting mailboxes and lines waiting mailboxes, and the maximum number of unanswered interrupts and requests waiting for a mailbox recorded during the time interval. A line is produced for each datanet (see figure 9-16.4). NOTE: This report is no longer produced.

SECTION 11. CENTRAL PROCESSING UNIT AND TAPE REDUCTION PROGRAM (CPUTRP)

11.1 Introduction

The Central Processing Unit and Tape Reduction Program (CPUTRP) is a FORTRAN program that sequentially processes the data recorded on tape by the CPU Monitor and the Tape Monitor of the General Monitor Collector (GMC). It produces a number of reports depicting the usage of the CPUs or the usage of the tape subsystem during the monitoring period. A list of these reports is found in table 11-1 and report descriptions are presented in subsection 11.5.

There are two inputs to the CPUTRP. The first is the data tape produced by the CPU Monitor and/or Tape Monitor in the General Monitor Collector. The second input is a set of report option control cards used to alter the reports in some way other than standard default. The various user input options and their formats are described in subsection 11.6. The actual reports produced by the CPUTRP are described in subsection 11.5.

11.2 Data Collection Methodology

The CPUM in the General Monitor Collector will create its own direct transfer traces (63 and 70) in order to collect data sufficient to analyze the utilization of the Central Processing Units. The method for generating these direct transfer traces is described in subsection 5.2.3 and the formats for the CPUM generated records used by the CPUTRP are described in subsection 5.4.4.

The Tape Monitor in the General Monitor Collector processes GCOS trace types 50, 51 and 52, and collects information to monitor the usage of the tape subsystem. The information collected on the occurrence of the above traces enables the CPUTRP to identify the jobs using tapes, which drives are used, how long a job is delayed due to nonavailability of tape drives, and the length of time a job is allocated to a given drive.

11.3 Analytical Methodology

There is no special analytical procedures used in this program. The program merely reports on the usage of the CPU and tape subsystem as it is reported by the General Monitor Collector.

11.4 Data Reduction Methodology

The CPUTRP is the only reduction program that actually produces reports from the data gathered by two different monitors (CPU and Tape Monitors). Therefore, a procedure needed to be developed whereby the user could specify which set of reports he desired, or if desired, allow him to produce the set of reports from each monitor. This capability could result in a given set of tapes needing to be analyzed as many as two times

within a given run. In order to manage this dual monitor reporting capability, a more complex user interface was required, then previously used by the other data reduction programs. This interface will be described in subsection 11.6.

11.5 CPUTRP Output

Output is divided, conceptually, into three parts - Execution and Error Reports, CPU Reduction Reports, and Tape Reduction Reports. These categories are shown in table 11-1, and further described below. An example of each report is displayed, and a simple explanation of how it was derived from the data is given.

11.5.1 Execution and Error Reports (Files 6, 7 and 8). These reports are written to three files: codes 6, 7 and 8. On file code 8, the user will find a printout containing processing or execution information - what the CPU and Tape Reduction Program found on the data tapes. Included in this information is the following:

- A list of the monitors in execution during the GMC data collection.
- An echo of the user input options, and the program's interpretation of them.
- The time, date, tape number and RCSR clock at the beginning of collection.
- If the timeframe option is used, a report of when the various timeframes were reached.
- A count of the traces reduced within each timeframe option.
- If a reduction period exceeded 9+ hours, an indication that a 35-bit internal clock overflowed (no error implied).
- The time, date, tape number and RSCR clock at end of data collection, but only if a termination trace is processed.
- When the NEW option is user selected, the items above are repeated on a new page.
- Version number - The software that is compatible with this documentation should indicate VERSION 7.2-11.74, 24 SEP 1982.

A sample of this report is shown in figure 11-1.

Errors, except for tape handler errors, are listed on file code 6. Any such occurrence is of concern to C751 and should be reported for

Table 11-1. Central Processing Unit and Tape Reduction Reports

CPU Reports:

a. A periodic tabular report which shows the cumulative CPU utilization and queuing by various program categories (default period is 10 minutes) (ID #0).

b. An integer valued histogram of the number of activities in the CPU queue (ID #3).

c. A real valued histogram of the CPU burst length for user activities (ID #1).

d. A real valued histogram of the CPU burst length for all activities (ID #2).

e. An integer valued histogram displaying the percentage of time a given activity spent waiting in the dispatcher's queue (ID #8).

f. A periodic plot of the CPU idle time (default period is 10 minutes) (ID #4).

g. A periodic tabular report of WIN CPU usage (default period is 10 minutes) (ID #5).

h. A report showing the CPU usage and queuing statistics for each activity processed (ID #6).

Tape Reports:

i. An integer valued histogram of the number of tape drives in use (time corrected) (ID #7).

j. A tabular report which describes, for each activity that used tapes, the device and channel utilization and the delay time waiting drives (no report ID # assigned).

Execution Reports:

k. A report which provides an overview of the data reduction - it describes the initial and final data tapes, the card input options and user selected time frames as they occur (no report ID # assigned).

l. A report which shows errors detected by the tape handler, and, if selected, debug output (no report ID # assigned).

m. A report which shows errors detected by program modules other than the tape handler (no report ID # assigned).

```
CPU AND TAPE REDUCTION, VERSION 7.2 - 11.74, 24 Sep 1982

     ACTIVE MONITORS:  ****  ****  CPUM  TM    CM    ****  ****  ****  ****  ****  ****  ****

     DATA COLLECTION BEGAN WITH TAPE D0020 AT 11:44:41 ON 82-12-15,
     RSCRA/Q WERE  000000000000 731542120046

CARD INPUT IS INTERPRETED AS FOLLOWS:

 1: HPRINT(1) CPRINT(300)
     HISTOGRAM REPETITION PERIOD IS 1          SECONDS MINIMUM.
     CPU REPORT REPETITION PERIOD IS NOMINALLY 300        SECONDS.

 2: NEW(D0020)
     TAPE D0020  WILL LEAD OFF A FOLLOWING REDUCTION.

      REDUCED          275   T70   TRACES.

     DATA COLLECTION ENDED WITH TAPE D0020 AT 11:56:48 ON 82-12-15
     RSCRA/Q WERE  000000000000 737066771640
```

Figure 11-1. Sample Execution Report

AD-A126 234 GENERALIZED MONITORING FACILITY USERS MANUAL CHANGE 4 (U) COMMAND AND CONTROL TECHNICAL CENTER WASHINGTON DC 81 JAN 83 CCTC-CSM-UM-246-82-CHG-4

UNCLASSIFIED F/G 9/2 NL

END

FILMED

DTIC

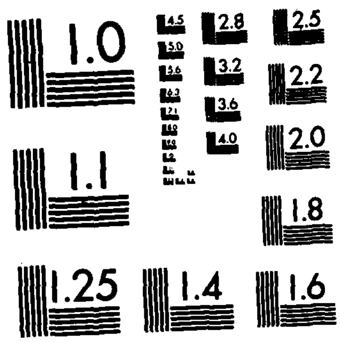

MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

interpretation and corrective action (202-695-0856). The occurrence of errors will, as a rule, invalidate the run

All tape handler messages, including tape error messages, are found on file code 7. The following messages are of informative value and are generally self explanatory.

- "MOUNTING ANOTHER REEL #..."

- "END REACHED IN NXTREC AT RCDCNT ..."
  Explanation: A tape or operator error forced the tape handler to treat the condition as an end-of-file. This message may be preceded by the following

- "ERROR ON THE TAPE READ"

- "TAPE MOUNTED APPEARS INCORRECT BUT IS NEW ... PROCESSING WILL CONTINUE ..."
  Explanation: A newly mounted continuation reel does not conform to continuation conventions, but is also not positively incorrect. The run is probably invalid.

- "INCORRECT TAPE MOUNTED 3 TIMES OR NEXT TAPE CANNOT BE READ. RUN ENDED".
  Explanation: A newly mounted continuation reel is positively incorrect. This will be preceded by

- "WRONG TAPE MOUNTED. WANTED ..."

- "JULIAN DATES DO NOT AGREE. RUN ENDED"
  Explanation: A new physical record bears an incorrect Julian date. This event may result from re-using an old GMC data tape in another collection run in which GMC termination was improper or incomplete.

- "GMFEXC FOUND END OF TAPE ON FIRST READ"

- "GMFEXC ERROR: FIRST RECORD IS IN ERROR OR WRONG TAPE MOUNTED. PHYSICAL RECORD DUMPED".
  Explanation: The occurrence of this message is of concern to C751 only if there was no tape mounting or tape number confusion. This may be preceded by

- "GMFEXC ERROR: ASKED FOR TAPE #... GOT ..."
  Explanation: The NEW option was used, and the user specified tape number does not agree with the number found on the NEW tape.

All other tape handler messages are of concern to C751 and should be reported for correction.

Debug output, when selected, shares file code 7 with the tape handler.

11.5.2 Central Processing Unit Monitor Reports. The CPU title page is printed to file code 14, immediately ahead of any histograms. The title page contains a summary of the systems configuration, the time reduction effectively started and stopped, as well as identifying the system which was monitored and the tape numbers containing the data set.

Following this information is a list of all dispatcher options that were active. This information is obtained from word 0 of the dispatcher module (.MDISP). If Priority B processing was enabled, then the SNUMBs being granted Priority B processing will be listed.

The user can alter the dispatcher algorithm to satisfy particular installation requirements. The three most commonly used algorithms are Urgency Thruput, I/O Priority and Priority B processing. If none of these options are selected, job urgency and channel time will be ignored and the standard rules for job selection will be observed. In effect, the standard rules imply a First-In/First-Out mode of operation.

While the dispatching algorithm involves some rather complex decisions, the primary driving force behind the algorithm is the urgency of a job. When the I/O Priority option is set, a job's urgency code is computed at 192 millisecond intervals of processor time. If the job is an ordinary slave program and its urgency code is not 0, the code is reduced by 1. Each time an urgency code is reduced to 0, a processor-bound job is dispatched at the next dispatch. If a job is a priveleged slave program, or if its urgency code is 0, the urgency code is recomputed as follows:

Urgency Code = L+T

where T = 1 if total channel time is at least equal to total processor time; otherwise T = 0.

L = 0,1,2,3,4 depending on the ratio of local channel time to local processor time.
set to 0 if ratio $<$ 1:1
set to 1 if ratio $\geq$ 1:1 $<$ 4:1
set to 2 if ratio $\geq$ 4:1 $<$ 16:1
set to 3 if ratio $\geq$ 16:1 $<$ 64:1
set to 4 if ratio $\geq$ 64:1

When the Urgency Thruput option is set, a job's dispatcher urgency is computed whenever the job's dispatcher urgency reaches 0. The dispatcher urgency is reduced by one each time the job is placed in the queue. A job's dispatcher urgency is computed as (U+2)/8 where U is a job's processing urgency from .CRSN1. This formula explains why it is very important to ensure that system programs and priority jobs are given processing urgency levels (U value) significantly higher than the ordinary slave job. Assume that $PALC has a processing urgency level of 54 and a

user job has a processing urgency level of 40. In this example, $PALC has a dispatcher urgency of (54+2)/8 = 7 while the slave job has an urgency level of (40+2)/8 = 5. Therefore, $PALC is given priority for only two dispatches until it will be considered equal to the slave job. On the other hand, if the slave job has a processing urgency of 6, then it would have a dispatching urgency equal to (6+2)/8 = 1. In this case, $PALC would be given priority for seven out of eight dispatches.

It is possible for a site to have both I/O Thruput and Urgency Thruput set. In this case, a combination of the above algorithms is implemented.

Finally, the Class B level of priority is a dispatching option which can be used to ensure faster processing for a set of one to three jobs. The order in which the jobs are given affects priority; that is, each job is checked for dispatch in the order in which it occurs. The Frequency Count is the number of dispatches to other jobs after two consecutive dispatches to this priority job. Thus, a Frequency Count of one allows the priority job two out of three dispatches. The Time Quantum is the number of 32-millisecond intervals allowed per time quantum. If both the Frequency and Time Quantum values are 0, priority for this job is delayed until a change priority request is issued.

The next several lines of output describe the overhead of all GMF monitors that were active during data collection. The monitor name is given, its CPU time in seconds, and its overhead as a function of total processor power. The GMF executive overhead is separated from the actual monitors and is listed as "EXEC". The monitor "NAME" is an area of code within the Mass Store Monitor and even though listed separately, it is also included under the monitor "MSM". The monitor "FMS" is also an area of code within the Mass Store Monitor, but in this case it has not been included under the monitor "MSM". Monitor "CM" describes the processor overhead of subroutine T4 (terminate processing) and subroutine T22 (start I/O processing). Monitor "MSM" describes the processor overhead of subroutine T7 (connect processing). Therefore, if the Channel Monitor was active, but the Mass Store Monitor was not, this report will list both "CM", "MSM" and "FMS" monitors. If both the Channel and Mass Store Monitors were active, then the combined overhead of both monitors can be found by the same sum above. For purposes of this report, percent overhead is computed as (CPU time used by monitor)/(total CPU time) x (number of processors).

If a termination record is not reduced for any reason, the lines describing monitor overhead will not be printed. Figure 11-2 provides a sample report.

11.5.2.1 <u>CPU Queue Length Distribution (File 14)</u>. Figure 11-3 shows a sample histogram of how busy the CPU queues were. It indicates how many activities were waiting for a processor at the time the samples were taken. This histogram will be produced only if the dispatcher queuing option of the CPU Monitor has not been disabled. It should be noted that

the dispatcher queue is not analyzed over a continuous time frame, but rather is sampled once every 1500 dispatches.

All histograms produced by the CPUTRP are interpreted in the following manner:

- In the INDIV NUMBER column, the histogram displays the number of occurrences of a particular event. The particular event being evaluated is represented by the figures in column 5. Therefore, in figure 11-3 we see that 231 times the CPU queue had a length of 0, while 44 times the queue length was 1.

- In the INDIV PROB. column, the histogram displays the probability that a given event will occur. Therefore, there is an 84% probability that the queue length was 0; i.e., 84% of the time there was no queue, while 16% of the time the queue length was 1.

- The CUMUL NUMBER and CUMUL PROB. columns are merely the cumulative probability distribution of the histogram.

```
CPU AND TAPE REDUCTION, VERSION 7.2 - 11.74, 24 SEP 1982

****************************************************************************************************
****************************************************************************************************
*****                                                                                         *****
*****                ***** T H E   C P U   A N D   T A P E   M O N I T O R *****              *****
*****                                                                                         *****
*****                                                                                         *****
*****                                                                                         *****
*****                     REDUCTION STARTED AT 1730:20.349 WED 82-12-27                       *****
*****          AND COMPLETED AT 1909:10.372 WED 82-12-27 A TOTAL TIME OF 1.65 HOURS           *****
*****                     ON SYSTEM NMCC2  RUNNING 4J2 ON TAPE D0020                          *****
*****                                                                                         *****
*****                  THE SYSTEMS CONFIGURATION CONSISTED OF:                                *****
*****                                                                                         *****
*****                            2 - 6680  CENTRAL PROCESSORS                                 *****
*****                                                                                         *****
*****                                 DISPATCHER OPTIONS                                      *****
*****                            IN CORE PUSH AREA                                            *****
*****                            DYNAMIC BUFFERING OF SSA MODS                                *****
*****                            CLASS B PRIORITY ON                                          *****
*****                                                                                         *****
*****                                   PRIORITY B JOBS                                       *****
*****                        SNUMB          FREQ       TIME QUANTUM                           *****
*****                         TS1             2             7                                 *****
*****                                                                                         *****
*****                                                                                         *****
*****             MONITOR OVERHEAD   -  MONITOR      TIME(SEC)     % OVERHEAD                 *****
*****                                   -------      ---------     ----------                 *****
*****                                    EXEC           47            3.26                    *****
*****                                    CPU            47            3.26                    *****
*****                                   -------      ---------     ----------                 *****
*****                                    TOTAL          94            6.52                    *****
*****                                                                                         *****
****************************************************************************************************
****************************************************************************************************
```

Figure 11-2. CPU Title Page

```
CPU AND TAPE REDUCTION, VERSION 7.2 - 11.74, 24 SEP 1982

                    DISTRIBUTION COLLECTE  ON SYSTEM NMCC2 AT 1730:20.349 WED 82-12-27

                                        CPU QUEUE LENGTH DISTRIBUTION

  INDIV.   CUMUL.   CUMUL.   INDIV.   QUEUE          PERCENT OF PROBABILITY OF OCCURRENCE              REPORT
  NUMBER   NUMBER    PRC      PRC     LENGTH  00  10  20  30  40  50  60  70  80  90  100                 1
                                              I....I....I....I....I....I....I....I....I....I....I
    231      231    84.00    84.00     0-   0 IXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
     44      275   100.00    16.00     1-   1 IXXXXXXXX

    275 ENTRIES TOTAL    AVERAGE =   0.16000    VARIANCE =   0.134    STANDARD DEVIATION =   0.367
```

Figure 11-3. CPU Queue Length Distribution

In figure 11-4, we find that 40% of all CPU bursts have a duration of 6ms or less while 96.5% of all CPU bursts have a duration of 12ms or less.

- The last line of the histogram provides some standard statistical output such as average, variance and standard deviations.

11.5.2.2 CPU Burst Distribution for User Activities/All Activities (File 14). Figures 11-4 and 11-5 show sample histograms of the CPU burst length distribution for user activities/all activities. They provide a measure of how long an activity held a processor before relinquish or interrupt.

WARNING: Each report is not a true histogram of individual burst lengths but is one of burst averages over a sampling period. For example, figure 11-5 shows 37062 bursts, but only 275 samples were taken. Each sample provided a count of intersample bursts and the accumulated burst time. For each sample, the histogram was charted as though each burst had the same length, namely, the average intersample burst length equals the accumulated-burst-time divided by count-of-bursts.

11.5.2.3 Percent of Memory Time in Queue. Figure 11-6 shows a sample histogram of the percent of time spent by an activity waiting in the dispatcher's queue for an available processor. During this wait time, an activity is effectively residing in memory, but accomplishing no useful work. This figure provides a quick answer to the question, "Is my system CPU-bound?" Any system in which activities are spending more than 15% of their memory time, waiting for a given resource, is definitely constrained by that resource. This histogram is only produced if the dispatcher queuing option of the CPU Monitor has not been disabled.

11.5.2.4 CPU Time Reports (File 10). This tabular report (figure 11-7) is produced every ten minutes (default period) of elapsed time. The first line indicates how many seconds have elapsed in the analysis, the total CPU time that has elapsed (total elapsed time x the average number of processors available through this period), the system ID, time of day and date. The amount of total CPU time is adjusted for any processors that have been released.

The next block of print gives the accumulated CPU time for each system program, for TSS, TRAX, WIN, and for all slave activities (presented as a single figure under the heading USER). The CPU time used by the slave portion of the GMC collector is shown separately.

```
CPU AND TAPE REDUCTION, VERSION 7.2 - 11.74, 24 SEP 1982

                         DISTRIBUTION COLLECTED ON SYSTEM NMCC2 AT 1730:20.349 WED 82-12-27

                                   CPU BURST DISTRIBUTION FOR USER ACTIVITIES ONLY

     INDIV.     CUMUL.     CUMUL.     INDIV.     MILLI-          PERCENT OF PROBABILITY OF OCCURRENCE            REPORT
     NUMBER     NUMBER      PRC        PRC       SECOND     00   05   10   15   20   25   30   35   40   45   50     3
                                                             I....I....I....I....I....I....I....I....I....I....I
       85         85        2.00       2.00       3.000      IXX
      144        229        5.30       3.30       4.500      IXXX
     1511       1740       40.40      35.10       6.000      IXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
      703       2443       56.70      16.30       7.500      IXXXXXXXXXXXXXXXX
     1593       4036       93.70      37.00       9.000      IXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
       75       4111       95.40       1.70      10.500      IXX
       47       4158       96.50       1.10      12.000      IX
      150       4308      100.00       3.50      13.500      IXXX

     4308 ENTRIES TOTAL      AVERAGE =    6.91398    VARIANCE =    3.543    STANDARD DEVIATION =    1.882
```

Figure 11-4. CPU Burst Length Distribution for User Activities

```
CPU AND TAPE REDUCTION, VERSION 7.2 - 11.74, 24 SEP 1982

                    DISTRIBUTION COLLECTED ON SYSTEM NMCC2 AT 1730:20.349 WED 82-12-27

                               CPU BURST DISTRIBUTION FOR ALL ACTIVITIES

   INDIV.    CUMUL.    CUMUL.    INDIV.    MILLI-       PERCENT OF PROBABILITY OF OCCURRENCE          REPORT
   NUMBER    NUMBER     PRC       PRC      SECOND   00   10   20   30   40   50   60   70   80   90   100       4
                                                     I....I....I....I....I....I....I....I....I....I....I
     707       707      1.90      1.90      1.500    IX
     574      1281      3.50      1.50      3.000    IX
   27974     29255     78.90     75.50      4.500    IXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
    3979     33234     89.70     10.70      6.000    IXXXXX
    1698     34932     94.30      4.60      7.500    IXX
    1834     36766     99.20      4.90      9.000    IXX
      94     36860     99.50      0.30     10.500    I
      48     36908     99.60      0.10     12.000    I
     152     37060    100.00      0.40     13.500    I
       1     37061    100.00      0.00     15.000    I
       1     37062    100.00      0.00     16.500    I

   37062 ENTRIES TOTAL     AVERAGE =     4.24950     VARIANCE =    2.104     STANDARD DEVIATION =    1.451
```

Figure 11-5. CPU Burst Length Distribution for All Activities

CPU AND TAPE REDUCTION, VERSION 7.2 - 11.74, 24 SEP 1982

DISTRIBUTION COLLECTED ON SYSTEM TS-7.2 AT 1730:20.349 WED 82-12-08. INITIAL GMC TAPE #00673

PERCENT OF MEMORY TIME IN QUEUE

INDIV. NUMBER	CUMUL. NUMBER	CUMUL. PRC	INDIV. PRC	%		PERCENT OF OCCURRENCE 00 05 10 15 20 25 30 35 40 45 50 I....I....I....I....I....I....I....I....I....I....I	REPORT 7
11	11	4.911	4.911	0-	1	IXXXXX	
4	15	6.696	1.786	2-	3	IXX	
0	15	6.696	0.	4-	5	I	
3	18	8.036	1.339	6-	7	IX	
6	24	10.714	2.679	8-	9	IXXX	
4	28	12.500	1.786	10-	11	IXX	
1	29	12.946	0.446	12-	13	I	
5	34	15.179	2.232	14-	15	IXX	
2	36	16.071	0.893	16-	17	IX	
5	41	18.304	2.232	18-	19	IXX	
2	43	19.196	0.893	20-	21	IX	
2	45	20.089	0.893	22-	23	IX	
3	48	21.429	1.339	24-	25	IX	
4	52	23.214	1.786	26-	27	IXX	
.	.	.	.	.	.	.	
.	.	.	.	.	.	.	
.	.	.	.	.	.	.	

Figure 11-6. Percent of Memory Time in Queue

CPU AND TAPE REDUCTION, VERSION 7.2 - TEST, 28 SEPTEMBER 1982
ELAPSED (PROCESSOR) TIME IS 1806.06 (3612.12) SECONDS ON SYSTEM NMCC2 AT 1241:40.061 MON 82-09-13
CPU TIME (EXCEPTING OVERHEAD AND IDLE) IN HUNDREDTHS SECOND, # OF DISPATCHES, SERVICE TIME IN CLOCK PULSES AND % CPU OF TOTAL-CPU

	CALC	PALC	SYOT	RTIN	T/DS	LOGN	FSYS	HEALS	GMC	TSS	NCP	TLNT	FTS	TLCF	TRAX	USER
CPU	493	1398	2752	221	0	367	0	0	13	23780	1499	186	1682	1	0	120361
DSP	1765	5072	4937	708	0	2499	0	0	67	100832	2983	498	3970	6	0	55355
STM	178	176	356	200	0	94	0	0	133	150	321	239	271	134	0	1391
CP%	0.14%	0.39%	0.76%	0.06%	0. %	0.10%	0. %	0. %	0.00%	6.58%	0.42%	0.05%	0.47%	0.00%	0. %	33.32%

TSS:	CPU IN HUNDREDTHS SECOND				NUMBER OF DISPATCHES				SERVICE TIME IN CLOCK PULSES			
	TS1	TS2	TS3	TS4	TS1	TS2	TS3	TS4	TS1	TS2	TS3	TS4
EXEC	11730	0	0	0	82132	0	0	0	91	0	0	0
SDSP	12049	0	0	0	18700	0	0	0	412	0	0	0

OVERHEAD, IDLE, GATE LOOP TIME IN HUNDREDTHS SECOND (GATE LOOP ALSO INCLUDED IN OVERHEAD)

	CPU 1	CPU 2	CPU 3	CPU 4	CPU 5	CPU 6	TOTAL
OHEAD	32796	10580	0	0	0	0	43377
IDLE	109875	56832	0	0	0	0	166708
GLOOP	557	657	0	0	0	0	1215

# CPUS,	% SYSTEM CPU,	% TSS CPU,	% WIN CPU,	% TRAX CPU,	% USER CPU,	% IDLE TIME &	% GATE LOOP OF BUSY...	SINCE
2.00	13.46	6.58	0.93	0.	33.32	46.15	0.62	RUN START
2.00	16.19	9.69	0.05	0.	46.70	27.94	0.58	LAST PRINT

Figure 11-7. CPU Time Report (Part 1 of 2)

```
QUEUE TIME IN TENTHS SECOND, AVERAGE QUEUE TIME PER DISPATCH IN CLOCK PULSES AND % QUEUE OF ELAPSED-TIME SINCE RUN START
     CALC    PALC    SYOT    RTIN    T/DS    LOGN    FSYS    HEALS     GMC      TSS    NCP    TLNT    FTS    TLCF    TRAX    USER
QUE   124     188      75      16       0    8177       0        0       2      472   4227      25   2484    3746       0   16477
AVG   450     238      97     145       0   20941       0        0     205       29   9069     327   4005 3996174       0    1905
QU% 0.69%   1.05%   0.42%   0.09%   0.  %  45.28%   0.  %    0.  %  01.01%    2.62% 23.41%   0.14% 13.76%  20.74%   0.  %
---AVERAGE QUEUE POSITION WHEN IN QUEUE (OR PROCESSOR)
AQP  2.00    1.33    1.29    2.00    0.      1.17    0.                       0.99   1.07    2.00   1.36    0.
QUEUE TIME IN TENTHS SECOND, AVERAGE QUEUE TIME PER DISPATCH IN CLOCK PULSES AND % QUEUE OF ELAPSED-TIME SINCE LAST PRINT
     CALC    PALC    SYOT    RTIN    T/DS    LOGN    FSYS    HEALS     GMC      TSS    NCP    TLNT    FTS    TLCF    TRAX    USER
QUE    58     103      12       6       0     768       0        0       1      235      5       2      3       0       0    7541
AVG   634     298     372     123       0   11525       0        0     528       31    409     77C    598     409       0    1610
QU% 0.96%   1.72%   0.21%   0.11%   0.  %  12.72%   0.  %    0.  %   0.02%    3.89%  0.10%   0.04%  0.05%   0.00%   0.  %
---AVERAGE QUEUE POSITION WHEN IN QUEUE (OR PROCESSOR)
AQP  2.33    1.44    1.00    2.00    0.      0.      0.                       0.96   0.      0.     0.      0.

TSS:            QUEUE IN TENTHS SECOND                                          AVG QUEUE TIME IN CLOCK PULSES
             TS1      TS2      TS3      TS4                                      TS1      TS2      TS3      TS4
EXEC         438        0        0        0                                       34        0        0        0
SDSP           0        0        0        0                                        0        0        0        0

 AVERAGE QUEUE LENGTH WAS  1.39 SINCE RUN START,  1.50 SINCE LAST PRINT.
```

Figure 11-7. (Part 2 of 2)

The second line of print in this block gives the accumulated number of dispatches for each of the associated programs. The service time for each program is presented on the third line. This figure is calculated by dividing the accumulated CPU time by the accumulated number of dispatches. This figure is then converted to clock pulses (1/64 ms). The final line in this block indicates what percentage of the total processor power is being used by each of the associated programs. This figure is based on 100% power availability. Therefore, if the operating system is using 30% processor power on a 3-processor system, it would be using (30%)(3) = 90% of a processor.

The third set of print provides, for each copy of Time Sharing (TSS), the CPU time attributable to the executive phase and to the subdispatch phase. In addition, for each copy of TSS, the average service time is also presented. This figure is represented in clock pulses (1/64ms). This figure represents the average amount of CPU time used by TSS whenever it received control of the processor. By tracking this figure, it is possible to see if bad periods of TSS response coincide with periods of time when TSS was receiving inadequate time slices of CPU service. The service time is calculated by taking the total CPU time accumulated by TSS and dividing it by the total number of CPU bursts accumulated by TSS.

If the user has selected the option to monitor special SNUMBs, the next block of print will list each of the special SNUMBs under which will appear the total CPU time accumulated by each of the SNUMBs. This block of print is not shown on the sample figure.

The next block of print shows the amount of overhead, idle and gate-loop time accumulated by each processor and by total. Gate-loop data is not applicable in a single processor environment. Gate loop time is that amount of time a processor is locked from executing because another processor has locked a required table or blocked a given area of code. In a multiprocessor environment, there are many instances where one processor is required to alter the values in a given table. While these values are being changed, the system wants to ensure that another processor does not reference the table, while it is in the midst of being changed. To prevent this from happening, the system will lock the table while it is in the midst of being altered. If a second processor desires to reference a locked table, it is required to execute a CPU "dead" loop while it waits for the table to be opened. The GLOOP statistics display the amount of such loop time executed by each processor. This value will not be reported for a single processor environment but should appear for any multiprocessor environment where the software release was WW7.2/4JS. If this data is not reported, it indicates a problem with the GMF collector and CCTC should be contacted.

The final block of print is a percentage breakdown of CPU usage into the categories SYSTEM, TSS, TRAX, WIN, USER, and IDLE; also shown are the percentage of gate-loop time relative to processor busy time and the time corrected number of processors in use. These figures are printed both for the current 10-minute interval and for the total current reduction

interval. This allows the user to track when peak usage of CPU power is occurring and what portion of the system is using this power.

Notes: (1) SYSTEM time includes CPU time accumulated in the "functions" OVERHEAD, CALC, PALC, SYOT, RTIN, TDS, LOGN, FSYS, DMTEX and MONITR. The time attributable to WIN, TSS and TRAX is neither SYSTEM nor USER. (2) The definition of overhead time is the time spent in the interrupt handler, the dispatcher and the SWAP processor, plus all gate-loop time. (3) It should be noted that the percentage figures are based on total CPU power and therefore add up to 100% (excluding the gate-loop percentage). In order to determine the "amount of a processor" required or used by a given function, it would be necessary to multiply the percentage figure for the function by the time corrected number of processors in use.

If the user has disabled the dispatcher queuing option of the CPU Monitor, the above blocks of output will occur two per page. If the dispatcher queuing option is enabled, then several blocks of output will be presented which depict the level of queuing occurring on the system.

The first block presents three lines of output for the same set of jobs as listed in the processor usage portion of the report. The total queue time accumulated by the job is given in line 1, the average queue time per dispatch is given in line 2 (presented in clock pulses) and the percent of elapsed time for which this job was in queue is presented in the third line. This final figure is not provided for the user entry. These three lines of output are accumulated data, since the start of the run. The next line of output indicates the average position this job held in the dispatcher's queue. If the job was being serviced by the processor, when the queue was examined, its relative queue position is considered to be zero.

These same four lines of output are repeated, but the data represents the various jobs' queuing statistics only since the last print out.

The final two blocks of output provide queuing statistics for the TSS and all special jobs selected for analysis. It should be noted that all TSS queuing statistics are for the executive only. Subdispatch queuing is not currently monitored.

11.5.2.5 CPU Plot of Percent Idle (File 31). This plot report shows the percentage of CPU power that is idle during each 10-minute (default period) interval. The data is taken directly from the CPU Time Report described in subsection 11.5.2.4. One horizontal line is output for every CPU Time Report table. The horizontal line represents one increment on the X-axis and it "paints" one datum of percent idle and the change of that datum since it was last plotted. Every 10th line also displays the current time of day. By the nature of the printing mechanism, an ordinate position is a cell, a range of values. The cell width is specified by a DELTA value given in the heading. The plot contains a heading line giving the system id, starting time of reduction and the date. It also contains the report title, and identifying mark, "A" here, of the curve, and the DELTA value. A plot summary, if not user deactivated, follows each plot.

Following the summary is the overall system idle percentage in three intervals; zero to 25% idle, 25 to 50% idle and 50 to 100?% idle. A final line shows the plot interval in seconds (default is 600 seconds). Figure 11-8 is a sample of this plot.

11.5.2.6 WIN Report. This report will be generated under the same time interval control as is used for the CPU Time Report (subsection 11.5.2.4) (default value every 10 minutes). Figure 11-9 is a sample of this report (no WIN software was active at the time this sample report was generated). For each WIN program, a single line of information is presented. This line indicates the total amount of CPU time accumulated by the program during the specified time interval, the number of CPU bursts accumulated during the time interval and the rate (bursts/sec) at which the bursts were generated. This report can be used to track those periods of time when WIN programs were using excessive CPU time or, on the other hand, were being denied CPU service.

11.5.2.7 CPU Access by SNUMB Report. This report (figure 11-10) summarizes CPU usage and queuing statistics for every activity processed during the monitor session. The various columns of the report are self-explanatory and require no special explanation. If a set of asterisks appear, for a given job, under the Average Queue Position column, this is an indication that this job was never caught in the processor queue. It should be stressed that the processor queue is not analyzed continually. Rather, the queue is examined under a sampling scheme and, therefore, it is possible that a given job will never be caught in the processor's queue.

11.5.3 *Tape Monitor Reduction Reports*. The tape reduction title page is identical to the CPU reduction title page (see subsection 11.5.2), except that the configuration described is pertinent to the tape subsystem. Shown are the channel number, the IOM number, the number of drives configured to the channel, the type of drive and whether the drives are cross-barred. The data shown is the configuration as presented in the boot deck. If any drives have been taken off line for maintenance or repair, it will not be reflected. The title page precedes any histograms; an example is presented as figure 11-11. It is written to file 14.

11.5.3.1 Number of Tape Drives in Use (Time Corrected) (File 14). A histogram report, seen as figure 11-12, shows the number of tape drives in use at the sampling epoch. The data are corrected for the inter-sample period, so that the figures listed under the heading "INDIV. PROB." correctly represent the fraction of reduction time the corresponding "NUMBER (of) DRIVES" were in use.

11.5.3.2 Tape Activity Report (File 13). This tabular report is made for each activity of a job that used tapes. For each activity of a job, the tapes used by that activity are described by type, unit number, and channel number. The report also presents how long the activity was

```
CPU AND TAPE REDUCTION, VERSION 7.2 - 11.74, 24 SEP 1982
                          DISTRIBUTION COLLECTION ON SYSTEM NMCC2 AT 11:44:41 ON 82-12-15

                                          CPU UTILIZATION OF SYSTEM/USER IDLE
PLOT-1

 DELTA=    7.9365E-01    A=IDLE

TIME OF DAY                                                 % IDLE

Y---------Y---------Y---------Y---------Y---------Y---------Y---------Y---------Y---------Y---------Y---------Y-------
--Y

X 11:45:11.14                                                                                                     A
I                                                                                                                AA
I                                                                                                       AAAAAAAAAA
I                                                                                                       AAAAAAA
I                                                                                     AAAAAAAAAAAAAAAAAAAAAAAAA
I                                                                                     A
I                                                                                    AA
I                                                                                    AA
I                                                                                    AA
X 11:49:53.28                                                                        AAA
I                                                                                     AA
I                                                                                  AAA
I                                                                                  AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
I
AAAAA

     CURVE          MINIMUM           MAXIMUM
     IDLE           0.6600E 02        0.9300E 02

                         0.  %  25% IDLE       0.  %  = 25% AND  = 50% IDLE      100.00%   50% IDLE

                        THE PRINT PERIOD FOR THIS PLOT IS NOMINALLY          30 SECONDS
```

Figure 11-8. CPU Plot of Percent Idle

DISTRIBUTION COLLECTED ON SYSTEM DSCC1 AT 1005:54.955 TUE 82-01-19. INITIAL GMF TAPE # 1182

WIN REPORT -- TIME IN HUNDREDTHS SECONDS; RATE, 1/SEC

	NCP			TELNET			FTS			TLCF		
TIME	CPUTIME	BURSTS	BURST-RATE	CPUTIME	BURSTS	BURST-RATE	CPUTIME	BURSTS	BURST-RATE	CPUTIME	BURSTS	BURST-RATE
1015:54.385	0	0	0.	0	0	0.	0	0	0.	0	0	0.
1025:56.385	0	0	0.	0	0	0.	0	0	0.	0	0	0.
1035:58.074	0	0	0.	0	0	0.	0	0	0.	0	0	0.
1045:54.581	0	0	0.	0	0	0.	0	0	0.	0	0	0.
1055:59.947	0	0	0.	0	0	0.	0	0	0.	0	0	0.
1106:01.044	0	0	0.	0	0	0.	0	0	0.	0	0	0.
1116:03.416	0	0	0.	0	0	0.	0	0	0.	0	0	0.
1121:59.169	0	0	0.	0	0	0.	0	0	0.	0	0	0.

Figure 11-9. WIN Report

CPU AND TAPE REDUCTION. VERSION 7.2 - 11.74. 24 SEP 1982

DISTRIBUTION COLLECTED ON SYSTEM TS-7.2 AT 1730:20.349 WED 82-12-08. INITIAL GMC TAPE #00673

CPU ACCESS BY SNUMB

TIME OF TERMINATION	JOB SNUMB	CPU TIME IN SECONDS	QUEUE TIME IN SECONDS	AVERAGE CPU QUEUE POSITION	MEMORY TIME (SEC)	PERCENT TIME IN QUEUE	ACTIVITY NUMBER	SWAP TIME (SEC)
1731:34	LCDTS	3.104E-01	2.704E 00	4.3	4.000E 00	67	1	0.
1732:04	8247T	1.774E 01	5.883E 01	3.2	1.030E 02	57	1	0.
1732:16	8278Z	1.686E 01	3.630E 01	3.8	7.700E 01	47	1	0.
1733:31	8287T	3.230E 01	1.013E 02	3.5	1.890E 02	53	1	0.
1734:09	8311T	6.848E-01	6.763E 00	4.6	1.000E 01	67	1	0.
1734:09	8310T	8.323E 00	1.675E 01	2.8	4.100E 01	40	1	0.
1735:22	8210T	3.156E 00	4.325E 01	2.1	5.700E 01	75	1	0.
1736:34	8318T	1.585E 00	9.959E 00	4.1	2.100E 01	47	1	0.

Figure 11-10. CPU Access by SNUMB

CPU AND TAPE REDUCTION, VERSION 7.2 - 11.74, 24 SEP 1982

***** T H E C P U A N D T A P E M O N I T O R *****

MONITORING ON 82-12-15 STARTED AT 11:44:41 AND COMPLETED AT 11:56:47 FOR A TOTAL TIME OF 0.20 HOURS
ON SYSTEM NMCC2 RUNNING W64000 ON TAPE D0020

THE SYSTEMS CONFIGURATION CONSISTED OF:

CHANNEL	IOM	# DEVICES	TYPE	SCT ADDR(SAME=XBAR)
18	0	6	9 TRACK TYPE	1936
19	0	6	9 TRACK TYPE	1960
24	1	6	9 TRACK TYPE	2032

MONITOR OVERHEAD -

MONITOR	TIME(SEC)	% OVERHEAD
EXEC	47	3.26
TAPE	1	0.01
TOTAL	48	3.27

Figure 11-11. Tape Reduction Title Page

```
CPU AND TAPE REDUCTION, VERSION 7.2 - 11.74, 24 SEP 1982

                    DISTRIBUTION COLLECTED ON SYSTEM NMCC2 AT 11:44:41 ON 82-12-15

                       THE TIME CORRECTED NUMBER OF TAPE DRIVES IN USE

  INDIV.   CUMUL.   CUMUL.   INDIV.   NUMBER            PERCENT OF PROBABILITY OF OCCURRENCE          REPORT
  NUMBER   NUMBER   PROB.    PROB.    DRIVES   00   10   20   30   40   50   60   70   80   90  100     10
                                               I....I....I....I....I....I....I....I....I....I....I
 4769030  4769030   1.000    1.000     1-   1  IXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX

 4769030 ENTRIES TOTAL       AVERAGE =    1.00000     VARIANCE =    0.     STANDARD DEVIATION =    0.
```

Figure 11-12. Number of Tape Drives in Use (Time Corrected)

delayed in peripheral allocation and how many drives it was waiting for. The final print tells how many drives were in use when the activity started and how many were in use when it ended. At the right margin of the report are two columns of numbers. The first column is the numerical sequence in which the activities started, and the second column is the numerical sequence in which they terminated. Refer to figure 11-13 for a sample report.

11.5.3.3 Tape Status Report (File 12). This tabular report may be seen as figure 11-14. It shows the status of all configured tape drives, and is repeated each time an activity was delayed due to unavailability of tape drives. (The first report entry is listed with the initial trace captured by the monitor collector). The reason for the delay, the epoch of the delay and the program number of the activity delayed are printed. As part of the status information, the following is presented: IOM-channel number, unit number, released or not, assigned or not, dedicated or not, used for T&D's or not, number of errors on the device, and the SNUMB of the job using that drive.

11.6 Default Option Alteration

The CPU and Tape Reduction Program uses two inputs. The first is the data tape(s) produced by the GMC. The second input is a set of report option control cards. The various options and their formats are described in the following subsections. Report option control cards are not required if the default options are acceptable to the user.

11.6.1 General Format. The general format for an option request is as follows:

```
COMMAND (PARAMETER-LIST)
```

where COMMAND is a keyword specifying what action is to be taken and (PARAMETER-LIST), if required, provides data to accomplish the action. The COMMAND is interpreted through its sixth character, so the input must match the first six characters of the valid commands described below. Shorter commands must match exactly.

11.6.1.1 Valid Commands. Eleven commands are recognized by the card input interpreter. The chart below shows each command, the action to be taken, and in parentheses, the default inplications.

o CPRINT	Print the CPU Time Report and one datum of the CPU Idle Time Plot (items A and f and g of table 11-1) at a specified repetition period. (Print every 600 seconds).
o DEBUG	Activate the debug statements in one or more subroutines. (All debug statements are inactive).

```
CPU AND TAPE REDUCTION, VERSION 7.2 - 11.74, 24 SEP 1982
   TAPE ACTIVITY REPORT FOR SYSTEM NMCC2 ON 82-12-15 AT 11:44

 TAPES FOR PROG (22)  MUM   ACTIVITY   1          TYPE          UNIT             CHANNEL NUMBER
                                                9 TRACK TAPE      4                  18( 0)
        THIS ACTIVITY HAD NO WAIT TIME


   WHEN ACTIVITY STARTED    1  TAPES WERE IN USE. WHEN IT ENDED     0  TAPES WERE IN USE          1      3
                              ***THIS PROGRAM WAS ACTIVE WHEN MONITOR STARTED
****************************************************************************************************************
```

Figure 11-13. Tape Activity Report

CPU AND TAPE REDUCTION, VERSION 7.2 - 11.74, 12 MAR 1982

TAPE STATUS REPORT FOR SYSTEM NMCC2 ON 82-12-15 AT 11:44

CHANNEL	UNIT	RLSED	ASGND	DED	T&D	ERRORS	SNUMB	CHANNEL	UNIT	RLSED	ASGND	DED	T&D	ERRORS	SNUMB
0-18	0	NO	NO	NO	NO	0		0-18	1	YES	NO	NO	NO	0	
0-18	2	NO	NO	NO	NO	0		0-18	3	NO	NO	NO	NO	0	
0-18	4	NO	YES	NO	NO	27	MUM	0-18	5	NO	NO	NO	NO	0	
0-19	0	YES	NO	NO	NO	0		0-19	1	YES	NO	NO	NO	0	
0-19	2	YES	NO	NO	NO	0		0-19	3	YES	NO	NO	NO	0	
0-19	4	YES	NO	NO	NO	0		0-19	5	YES	NO	NO	NO	0	
1-24	0	YES	NO	NO	NO	0		1-24	1	YES	NO	NO	NO	0	
1-24	2	YES	NO	NO	NO	0		1-24	3	YES	NO	NO	NO	0	
1-24	4	YES	NO	NO	NO	0		1-24	5	YES	NO	NO	NO	0	

****

****PROGRAM NUMBER 18 WAS DELAYED AT 11:53 FOR 2 7 TRACK TAPE

CHANNEL	UNIT	RLSED	ASGND	DED	T&D	ERRORS	SNUMB	CHANNEL	UNIT	RLSED	ASGND	DED	T&D	ERRORS	SNUMB
0-18	0	NO	NO	NO	NO	0		0-18	1	YES	NO	NO	NO	0	
0-18	2	NO	NO	NO	NO	0		0-18	3	NO	NO	NO	NO	0	
0-18	4	NO	YES	NO	NO	34	MUM	0-18	5	NO	NO	NO	NO	0	
0-19	0	YES	NO	NO	NO	0		0-19	1	YES	NO	NO	NO	0	
0-19	2	YES	NO	NO	NO	0		0-19	3	YES	NO	NO	NO	0	
0-19	4	YES	NO	NO	NO	0		0-19	5	YES	NO	NO	NO	0	
1-24	0	YES	NO	NO	NO	0		1-24	1	YES	NO	NO	NO	0	
1-24	2	YES	NO	NO	NO	0		1-24	3	YES	NO	NO	NO	0	
1-24	4	YES	NO	NO	NO	0		1-24	5	YES	NO	NO	NO	0	

****

Figure 11-14. Tape Status Report

- o HISTOGRAM — Modify a histogram report. (Histogram report parameters are as shown in table 11-2).

- o NEW — Prepare to accept another, or repeat, set of GMC data tapes and another set of option alteration commands for a following reduction run. (Reduction is complete when the current tape(s) are reduced).

- o HPRINT — Print the histograms only if the reduction interval is, at least, a specified minimum. (Print the histograms only if the reduction interval is at least 900 seconds).

- o ON OFF —
  (1) Turn on/off one or more reports. (All reports are on).

  (2) Turn on/off the CPU reduction module or the TAPE reduction module but not both. (CPU reduction is on; TAPE reduction, off).

- o PLOT — Modify a plot. (Plot report parameters are as shown in table 11-3).

- o STOP — Stop reduction after processing a specified number of physical tape records. (No limit).

- o TIMEFRAME —
  (1) Accept one to five time windows for overall data reduction or for plot output. (Data reduction and report output are not time limited).

  (2) Accept one to five time windows for debug output. (Debug output is not time limited).

  (3) Treat the receipt of a specified or implied special record as the end of one timeframe and start of another - that is, print the reports and start a new reduction period immediately. (No timeframe action is taken).

- o TITLE — Accept a new title to be printed on all reports as the new banner. ("CPU AND TAPE REDUCTION, VERSION 7.2-11.74, 24 SEP 82" is printed as banner).

11.6.1.2 Parameter List. The parameter list is a sequence of data items separated one from the next by a field of spaces. Any one space in the field may optionally be replaced by a comma or colon. A data item may be an alphanumeric, a quoted string, an integer, a real number or a null. Strings are enclosed in single or double quotes. Real numbers include the standard FORTRAN forms: for example, 12.34, 1234E-2, 1.234+1 are all

Table 11-2. Histogram Default Parameters

REPORT ID	TRAILER FLAG	# INTERVALS	LOW VALUE	INTERVAL SIZE
1	ON	40	0	1
2	ON	40	0	1
3	ON	45	0.0	1.5
4	ON	45	0.0	1.5
5	ON	40	0	1
6	ON	40	0	1
7	ON	40	0	1
8	ON	40	0	1
9	ON	40	0	1
11	ON	40	0	1

Table 11-3. Plot Default Parameters

REPORT ID	# POINTS to PLOT	Y-MINIMUM	Y-MAXIMUM
10	UNLIMITED	0.0	100.0

accepted as the same number. A null is two consecutive commas or colons; it stands for a missing data item and for the separator fields on either side of the missing item.

11.6.1.3 Command Syntax. Each command with its parameter list must be complete on a single card, but more than one command may be placed on each card if spaces separate them. The parameter list may be contiguous with the command keyword or spaces may intervene. The syntax for each command is given in following subsections. The data acronym RID stands for a one or two digit integer specifying the report identification number of the report to be affected. Report identification numbers are shown in table 11-1. The specifications "integer-W", "alphanumeric-X", "real-Y", and "string-Z" require user input of an integer number in a field at most "W" wide, an alphanumeric in a field at most "X" wide, a real number in a field at most "Y" wide and a quoted string in a field at most "Z" wide, respectively. Alphanumerics and strings exceeding the field width are accepted but are truncated to the right. Integers and reals exceeding the specified field width are rejected as illegal data. Integers may not contain a decimal point. Reals must contain a decimal point or an exponent. Reals are limited, for all data input, to a maximum of 14 characters including leading sign, decimal point, and exponent. Reals without exponent are further limited to a maximum of 11 characters including leading sign and decimal point. Data items shorter than the specified field width are entirely acceptable in all commands, but a null item is acceptable only where provided for.

11.6.1.4 General Control Command Syntax. The CPU and Tape Reduction program permits reduction of either the CPU associated traces or the tape subsystem associated traces during one pass over the data tapes. The CPU reduction phase is on, and the tape reduction phase is off, by default. The following commands may be used to exercise general control over the reduction:

- DEBUG — Activates all debug statements.

- DEBUG (MODULE-NAME-1,..., MODULE-NAME-N)
  Activates the debug statements in the explicitly named modules, where MODULE-NAME-i is the FORTRAN name given to the subroutine or function. Only the first six characters are interpreted. The routines currently debuggable are:

  CLOCK - keeps the master clock

  FIGCHG - examines a reconfiguration record for CPU related change

  GETOKE - lexical analysis of card input

LOGREC - gets a logical record meeting reduction requirements

NXTRECRD - tape input handler

RECNFG - gets the initial reconfiguration trace

TIMSET - sets up the array of start-stop times for user specified reduction windows

The one-shot program GMFEXC which reads the initial tape record is executed before card input. It may be debugged by setting program switch word, bit 11 (i.e., ON6).

- HISTOGRAM (RID, TRAILER-FLAG, NUMBER-OF-INTERVALS, LOW-VALUE, INTERVAL-SIZE)
  Turns on the histogram with report identification number RID, integer-2. Print/don't print the histogram summary if TRAILER-FLAG is ON/OFF, alphanumeric-3. The new number of histogram intervals is NUMBER-OF-INTERVALS, integer-3. The new low value and interval size are LOW-VALUE and INTERVAL-SIZE both integer-10 or real-14 according to the type of histogram. (In the current version, the number of intervals may not exceed 50).

The following two options are as above, but without change to the current specification of those parameters not appearing in the list.

- HISTOGRAM (RID,,NUMBER-OF-INTERVALS, LOW-VALUE, INTERVAL-SIZE)
- HISTOGRAM (RIT, TRAILER-FLAG)

- NEW (TAPE-NUMBER)
  Flags the reduction program to stop reading card input options and to proceed with data reduction. It also informs the program that another data reduction is to follow, and that the following reduction is to lead off with the tape number specified by this option (any six characters). For that following reduction, the default options will prevail unless modified by option alteration commands following the NEW command. This command must be the last command on a card (following commands on the same card are lost).

- HPRINT (X) Prints histograms if reduction interval is at least X, integer-10, seconds. Otherwise, the histogram data are lost.

- OFF Turns off all histograms and plots (if any).

- o OFF (RID-1,...,RID-n)
  Turns off the reports specified by report identification numbers RID-1,...,RID-n.

For the remaining OFF options, the first two characters of the parameter are necessary and sufficient.

- o OFF (PLOTS) Turns off all plots (if any).

- o OFF (HISTOGRAMS)
  Turns off all histograms.

- o OFF (CPU) Turns off the CPU reduction phase of the program and turns on the TAPE reduction phase.

- o OFF (TAPE) Turns off the TAPE reduction phase and turns on the CPU phase. This is the default specification and is not likely to be used, but is included here for completeness.

- o ON ON options use the same parameters as those shown for the OFF command, with appropriate reversal of intent. All reports are on by default.

Note: CPU and tape reduction cannot be done simultaneously. To accomplish both reductions in a single job, use the NEW option and a form of the ON/OFF option: ON(TAPE), for example.

- o STOP (X) Stop data reduction after reading X, integer-10, records maximum.

- o TIMEFRAME (DEBUG, START TIME, STOP TIME, START TIME, STOP TIME, ...)
  Accepts up to five start/stop pairs, where each component of the pair is a pseudo decimal number (real 7) in the form HHMM.SS where HH is the hour, MM is the minute and SS is the second. The decimal point must appear on the card, to further delimit the action of debug statements. The limiting action becomes effective as soon as the program determines "what time it is." This command must be accompanied by some form of the DEBUG command or no debug output will be produced. The timeframes are independent of, though not unaffected by, any other user selected timeframes.

- o TITLE (X) Accepts the string X, string-72, as the new banner for all reports. X must begin and end with a quotation mark, single or double, and may contain any FORTRAN acceptable character except the matching quote with which it is delimited.

11.6.1.5 CPU Reduction Command Syntax. In addition to those commands given in section 11.6.1.4, the following are pertinent to the CPU reduction phase.

o CPRINT (X) Prints the CPU Time Report and one datum of the CPU percent-idle plot every X, integer-10, seconds. The specified period is nominal and the actual period will vary with trace times. This command will not affect the on/off status of these reports.

o PLOT (RID, NUMBER-OF-POINTS, Y-MINIMUM, Y-MAXIMUM)
Turns on plot with report identification number RID, integer-2. The new number of points to plot is NUMBER-OF-POINTS, integer-10, maximum. The y-axis of plot is bounded below by Y-MINIMUM, real-14, and above by Y-MAXIMUM, real-14.

The following two options are as above, but without change to the value of those parameters not specified.

o PLOT (RID, NUMBER-OF-POINTS)
o PLOT (RID,, Y-MINIMUM, Y-MAXIMUM)

o TIMEFRAME (RID, START TIME, STOP TIME, START TIME, STOP TIME, ...)
Turns on report with report identification number RID, integer-2. Sets start and stop times according to the pair
START TIME, STOP TIME, START TIME, STOP TIME,
where each component is a pseudo decimal number (real 7) in the form HHMM.SS where HH is the hour, MM is the minute and SS is the second. The decimal point must appear on the card. Up to five start/stop pairs may be specified for a given report. The count of five may be achieved in one or several commands; however, the times must be in increasing order in order to get the results wanted. The final quadruple in a command may be abbreviated - if only a start time is given, the stop time is permanently unbounded. If RID is zero, the timeframe(s) bound the entire reduction but the status of individual reports (on or off) is unaffected. Histograms may not be individually delimited with this command, thus the command affects the entire reduction (RID equals zero) or the CPU plot (RID equals 10).

o TIMEFRAME (LOSTDATA)
Treat the receipt of a GMC lost data trace as the end of one timeframe and the start of another.

- o TIMEFRAME — Treat the receipt of a reconfiguration trace in which any CPU related data has changed (processor has been released or assigned) as the end of one timeframe and the start of another.

11.6.1.6 <u>Tape Reduction Command Syntax</u>. In addition to those commands given in subsection 11.6.1.3, the following is pertinent to the tape reduction phase.

- o TIMEFRAME (0,0,0, STOP TIME)
  By the nature of its construction, the Tape Reduction Module can only produce a nonwindowed set of reports (items M and N of table 11-1). A timeframe set for overall reduction (RID equals zero) will affect only the final stop time for these reports. The start time will always coincide with the start of the GMC collector. For consistency, the histogram of tape drive use (item L of table 11-1) will cover the same period. The stop time must be in the same format as described in previous timeframe commands.

11.6.1.7 <u>Card Input Errors</u>. If the card interpreter cannot understand a command keyword or data item, it will echo the entire card and mark the last correctly interpreted character position with the message "ILLEGAL COMMAND (or DATA) FOLLOWING...." Depending on the state of the interpreter, it may attempt to find the end of the incorrect command in order to pass on to a following command on the same card, or it may reject the remainder of the card; it will report which of the two actions was taken, and will then continue with card interpretation. In no case will an action confirmed by the interpreter be reversed or "undone" by a following error. Due to the manner in which cards are interpreted, the JCL file used to execute the CPUTRP should have the line numbers removed prior to execution. If this is not done, the interpreter will try to process the line numbers on the data cards as valid input requests and generate error messages when unable to do so. Despite the error messages, the CPUTRP will process correctly.

11.6.1.8 <u>Examples of Command Use</u>. A few examples will show the simplicity of the scheme.

- o ON (99 1 5 2, 3, 6)
  Action: Turns on reports 1, 2, 3, 5 and 6. Report 99 doesn't exist, so that parameter is ignored.

- o ON ("1", 002, THREE, 4.0)
  Action: Card is ignored since report id's are not one or two digit integers.

o TIMEFR (0 2102.00, 2304.00, 0500.00)
Action: Sets overall data reduction window to start at 2102 hours, stop at 2304 hours, and start again at 0500 hours with a permanently unbounded stop time.

o HISTOG (12, OFF, 42, 1.2, 1+4)
Action: If report ID #12 is a histogram, turns the trailer print flag off. If histogram #12 is of the real continuous type, turns the report on, sets the number of intervals to 42, beginning at 1.2 (Units) with each interval 10,000.00 (Units) wide.

o NEW (*A2345) DEBUG OFF (PL)
Action: Tape number *A2345 is remembered to be the lead off tape for a following run. Both DEBUG and OFF(PL) commands are lost since they appear on the same card. If these commands are desired for the next processing, they should appear on a set of following cards.

o DEBUG DEBUG (TIMSET)
Action: Turns on all debug statements. DEBUG(TIMSET) is a valid command, is correctly interpreted, and is repetitious.

11.7 JCL

The CPU and Tape Reduction Program is controlled by the following JCL:

```
$    SNUMB     ...
$    IDENT     ...
$    USERID    ...$...
```

These are required by the installation.

```
$    SELECT    B29IDPXO/OBJECT/CPU-TAPE
$    LIMITS    99,30K,,20K
```

These two cards load the reduction program object code and begin execution.

The following card,

```
$    TAPEn     01,X1D,,data-tape-number
```

describes the tape (7- or 9- track) generated by the data capture procedure, and

```
$    DATA      05
```

is used to identify the data cards that follow and that specify the input options as described in subsection 11.6.1. When options are not required, no $ DATA card is required.

Figure 11-15 shows a sample deck setup.

11.8 Multireel Processing

If more than a single reel of data has been collected, a series of messages will be outputted to the console informing the operator that a new data reel is required. The following are the messages produced.

a. DISMOUNT REEL #XXXXX THEN MOUNT REEL NUMBER YYYYY ON ZZZZZ

   In this case, XXXXX is the old reel number, YYYYY is the new reel number, and ZZZZZ is the tape drive ID.

   If the operator fails to mount the new tape, the above message will be repeated three times, after which the program will terminate, and all reports produced.

b. IS TAPE XXXXX MOUNTED ON DRIVE ID YYYYY (Y/N)

   In this case, XXXXX is the tape number being requested for mounting and YYYYY is the tape drive ID.

   This message occurs when the data reduction program finds the wrong tape has been mounted (by comparing internally generated tape labels). If the operator answers N, the message in (c) below is produced. If the operator answers Y, the data reduction program will terminate and all reports will be produced. In this case, the data reduction program is unable to process the tape. Even though the operator is mounting the correct tape, the internal label on the new tape does not match that being requested by the old tape. The user should check the data collection session to insure that the operator did not respond with an incorrect tape number during multireel change.

   After entering the Y or N, the operator will need to hit the EOM key twice in order for the response to be transmitted.

c. WRONG REEL JUST MOUNTED, DISMOUNT AND MOUNT REEL XXXXX ON ZZZZZ

   In this case, XXXXX is the new reel number, and ZZZZZ is the tape drive ID.

d. CAN TAPE XXXXX BE MOUNTED ON DRIVE YYYYY (Y/N)

   In this case, XXXXX is the new desired reel and YYYYY is the tape drive ID.

```
$        IDENT       18020251/30/3044
$        MSG2        THIS JOB WILL USE THE FOLLOWING TAPES
$        MSG2        IN THE GIVEN ORDER, ALL RING-OUT
$        MSG2        AJ283, AD281, D0020
$        SELECT      B29IDPXO/OBJECT/CPU-TAPE
$        LIMITS      99,30K,,20K
$        TAPE        01,X1D,,AJ283,,DATA
$        DATA        05
      DATA cards
$        ENDJOB
```

Figure 11-15. Sample Deck Setup

If the operator fails to answer this message it will be repeated until he responds with a "Y" for YES or "N" for NO. If he types in "Y", then message (a) will be repeated. If he types in "N", then the program will be terminated and all reports will be produced.

11.9 Tape Error Aborts

During the course of processing it is possible that the operator will be required to abort the data reduction program due to an irrecoverable tape error. If such a condition occurs, the operator should abort the job with a "U" abort. This will allow the data reduction program to enter its wrap-up code processing and produce all reports generated prior to the tape error.

THIS PAGE LEFT INTENTIONALLY BLANK

14.6.3.1 Obtaining the Data. The first step is for the user to examine the monitor data collection reports for the total time of the run. Any monitoring session of less than 2 hours in duration should be discarded. The Title Page will indicate the overhead generated by each of the GMF monitors active during the data collection phase. While this information is not used in this analysis, it is an item of information that is usually requested. Columns 1 and 2 of figure 14-4 can be obtained directly from the MUM Title Page. Under the System Configuration section of the Title Page, the amount of memory configured on the system will appear. This figure should be tentatively noted under column 5. Following this figure, the Title Page will list the amount of memory used by Hard Core, Core Allocator, SSA Cache and if any memory releases occurred during the monitoring session. All these functions have the effect of reducing the available user memory. These values should be summed and tentatively recorded under column 6.

Following the above information are several lines of statistics concerning the processing characteristics of the system. Columns 17 and 18 may be filled from this information. When determining the number of activities processed per hour, the analyst has two figures available. The analyst may choose to record the total number of activities processed per hour or he may record the number of (non-system schedular) activities processed per hour. During the course of the day many system schedular activities (activity 0 of any batch job) are processed. These activities are not really user generated, but rather are generated from the system. Therefore, by removing these activities from the total activities processed, a more realistic figure will be generated. The final three lines of the title page can be used to provide the data for columns 3 and 7 on the chart. Column 4 is filled from Report 1, columns 8 and 9 from Reports 11 and 12, columns 10 and 11 from Reports 16 and 17, column 12 from Report 19, column 13 from Report 51, column 14 from Report 37, column 15 from Report 50, and column 16 from Report 36. The two remaining columns that need to be completed are columns 5 and 6. The System Program Usage of Memory Report should be used to complete column 6. When processing the MUM data reduction program, the user should seriously consider using the MASTER input option. This will provide the user with a much better indication of the true system program load. In order to complete column 6, the user should record the total value appearing under the "WEIGHTED TTM" column of the System Program Usage of Memory Report. This value should then be added to the value already recorded under column 6. Finally column 5 can be determined by subtracting the value reported in Column 6 from the value previously reported in column 5.

14.6.3.2 Evaluating the Data. Figure 14-4 should be used in the following manner to determine if memory is a system constraint.

Step 1 - If column 7 shows a surplus of memory greater than 15% of the total available memory or greater than two times the value reported in column 4, then the implication is that there is no memory constraint on the system.

#1	#2	#3	#4	#5	#6	#7
Date	Hours	CPU/IO Ratio	Avg Activity Size (K)	Amt of Memory available for User activities	Amt of Memory used by System Functions	Memory Surplus or Memory Shortfall with and without the PALC Queues

#8	#9	#10	#11	#12
Avg Number of User Activities Waiting Memory	Avg Number of System Activities Waiting Memory	Avg Number of User Activities in Memory	Avg Number of System Activities in Memory	Ratio of Duration vs Memory Time

#13	#14	#15	#16	#17	#18
% Slave Memory Used	Amount of Time User Activity Swapped	Amount of Time Activity Waited for Original Memory Allocation	Number of User Swaps	Activities Per Hour (Throughput)	Total Swaps Per Hour

Figure 14-4. Memory Statistics

If Column 7 shows a surplus of memory, but does not exceed the aforementioned limits, the implication is that the current system has sufficient memory, but that the system is approaching memory saturation. If Colulmn 7 shows a shortfall of memory, and the value is greater than 15% of the total available memory or greater than two times the value reported in column 4, then the implication is that memory is a constraint on this system. Finally, if Column 7 shows a shortfall of memory, but does not exceed the aforementioned limits, the implication is that the system has reached memory saturation, but is still able to process the current workload.

Step 2 - It should be stressed that the value reported in column 7 is calculated over the entire measurement period and therefore could be biased by periods of heavy or light activity. It is for this reason that the user is urged to run the monitor during those periods of time where the workload is considered to be heavy and constant. In order to determine if the above type of biasing is occurring, the user may want to check Plots 1-3. If it appears that there is a mixing of light processing and heavy processing the user may want to re-run the data reduction program, using the time-frame option, to separate the heavy processing time.

Step 3 - Calculate the ratio of column 5 divided by column 4. This is an indication of the maximum number of user jobs that your system can support at any one time, without the occurrence of significant swapping. If the value in column 10 is equal to or exceeds this ratio, then the implication is that the system has reached memory saturation. If the value in Column 10 is within 2 units of the ratio, then the system probably has sufficient memory but is approaching saturation. Finally, if the value in column 10 is less than the ratio by more than 2 units, the current system has sufficient memory.

This step can be further verified by checking columns 14, 16 and 18 for indication of significant swapping.

Step 4 - If column 13 is less than 85, the current system should have sufficient memory and the other steps should not be showing indications of memory problems. If column 13 is between 85 and 95 then the current system is approaching saturation and at times may be showing some indications of a backlog. If the figure exceeds 95, then other steps should be indicating signs of moderate to severe memory problems.

Step 5 - If column 8 is greater than or equal to 3, the indication is that memory has become a constraining factor.

Step 6 - If Column 12 is greater than 2, the indication is that memory wait time is high and that memory is probably a constraining factor.

At this point, the user should have a fairly good indication as to whether or not memory is a constraining factor. The following steps will indicate some additional reports that the user should reference to determine those jobs that might be causing the memory problem.

Step 7 - One of the largest users of resources are jobs that abort and then must be rerun. Aborts usually occur due to user errors, but hardware aborts are not uncommon. If management is aware of aborting jobs and the reasons for them, they can possibly save substantial system resources. The Abort Report is described in subsection 6.3.8 and gives an indication of the system resources being wasted by aborting jobs.

Step 8 - It is important for management to be aware of jobs that are either misusing system resources or are requesting large amounts of system resources. Upon identifying such jobs, these jobs could be redesigned, scheduled for non-peak processing, or, in the case of wasted resources, the waste could be eliminated. The Excessive Resource Report allows the user to uncover jobs of this type and is described in subsection 6.3.10. When using this report, the following are suggested parameter values

wasted core - 10K
memory - either 35K (WWMCCS standard) or 2 times the value in column 4
CPU time - 15 minutes
IO time - 30 minutes
URG - 40
RATIO - 2

Step 9 - By examining the System Program Usage of Memory Report, the user can determine those system type jobs that are requiring memory. It is possible that some of the system jobs can be eliminated or at least reduced in size. This is especially true for the TSS. However, it must be realized that a limitation on Time Sharing size may adversely effect TSS response. In many cases, if large file transfers are being processed during prime time, the size of the FTS WIN subsystem can rise to 70 or 80K. By not allowing WIN file transfers to run during prime time, significant memory savings can result.

Step 10 - As is explained in great detail in subsection 6.3.15, it is vitally important that the overall urgency level of jobs being processed remain low. The Distribution of Urgency Report can be used to determine the overall urgency level of jobs being processed. This report should show that 60 percent of the jobs being processed at any one time have an urgency level below 40 and that a substantial proportion of these should have an urgency level between 5-10. The summary at the bottom should indicate that 75-80 percent of all activities processed had an urgency level below 20.

If this report indicates a large percentage of high-urgency jobs, then the SNUMB/IDENT report, or the Excessive Resource Report, can be used to identify those particular activities processing with a high urgency.

Step 11 - If the analyst wants to track the memory performance of a given set of jobs, the use of the SPECL input option and the generation of the Special Job Memory Reports will provide sufficient data for detailed memory tracking. This procedure is especially useful in analyzing the memory requirements of TS1, FTS and the special JDA-developed software (JDSIP, JDSUP). Refer to subsections 6.1.24 and 6.3.14 for complete descriptions of these Special Job Memory Reports.

Step 12 - Another indication of poor system performance possibly caused by memory shortfall, tape drive shortfall, poor operator performance or a poor system scheduler design is the long delay of jobs as they pass through the various allocation phases prior to core allocation. The Allocation Status Report, the System Scheduler Delay Time Histogram and the Delay Time Until Core Allocation Histogram can all be used to determine which jobs, and how many jobs, are being significantly delayed during the various allocation phases. These reports are all fully discussed in section 6.

Step 13 - Memory problems may also be occurring as a result of jobs being delayed due to CPU constraints or I/O constraints. In these cases, jobs tend to sit in memory due to a lack of other system resources. Because these jobs are being delayed, other jobs cannot enter memory, and memory demands begin to backlog. Therefore, if memory is a constraint, the user should consider conducting a CPU analysis as well as an I/O analysis.

14.6.4 CPU Evaluation. The CPU evaluation will determine the general utilization level of the processor and then determine if the CPU is dominated by GCOS or user program execution. In addition, the CPU evaluation can be used to determine if jobs are being significantly delayed by a lack of processor power. A CPU data reduction is required for this evaluation. It is also beneficial to have an associated MUM data reduction available for the same time period.

14.6.4.1 Data Recording. The heading page of the CPU data reduction report provides the dispatcher options currently in effect on the system. Recent tests have shown that the Urgency Thruput Option should be enabled, as well as the In-Core Push Area and Dynamic Buffering of SSA Modules. The TSS and the various WIN subsystems should not be placed in Priority B processing. In addition, sites should try to avoid enabling the I/O Thruput Option, unless strong justification exists to decide otherwise. The CPU Time Report is produced every 10 minutes of elapsed time and the data of interest should be found in the last 10-minute report.

THIS PAGE LEFT INTENTIONALLY BLANK

14.6.4.2 Evaluating the Data.

Step 1 - The CPU may by considered a bottleneck if the % Idle CPU column of the CPU Time Report is less than 20 or the summary column at the end of the CPU Idle plot indicates that 50% of the time or more, the processor is less than 25% idle. It is a fairly agreed upon standard that average processor utilization should not exceed 80%. By maintaining average processor utilization at this level, the processor will have sufficient remaining capacity to be able to handle those peaks of processor demand that normally occur during the day. The three category breakdowns at the end of the CPU plot represent the three conditions of insufficient power, sufficient power, and excess power respectively.

Step 2 - The % Gate Loop column of the CPU Time Report provides an indication of the percentage of CPU power being lost because multiple processors are interfering with each other. Within GCOS, there are many tables that must be updated and/or referenced by a processor during its execution. When these tables are being used, the processor must insure that they are not altered in any manner. In order to accomplish this, the processor will lock a gate. This locked gate will prevent any other processor from accessing this table. When the first processor has completed its use of the table, the gate will be opened. If a processor must access a gated table, it will simply perform a "CPU loop" at that table, waiting for it to be opened. The amount of time being spent in this gate locked-CPU loop code is depicted by this column. If this value is greater than 5%, then this is an indication that the multiple processor configuration is beginning to lose its cost effectiveness.

Step 3 - The MUM Excessive Resource Report can be used to determine those jobs requiring excessive CPU resources. In addition, the CPU Plot report can be used to determine those times of day when processor availability is in the critical range. Using these two items of information, system schedular classes can be created based on CPU requirements.

Step 4 - Another indication that the CPU is a bottleneck can be determined from the CPU Queue Length histogram. If the average queue length is greater than two times the number of processors configured, then the processors are being requested to handle an excessive amount of work. This queuing of jobs at the CPU is an indication that during some period of the day, there is insufficient CPU power to handle the workload. Once again, the CPU Plot can be used to determine those periods of time.

In addition, at the bottom of each 10-minute section of the CPU Time Report, a line is printed indicating the CPU queue length during the last 10-minute period, as well as since the start of the run. These figures provide an excellent indication of those times during the day that the processor is being overloaded. Control of the scheduler queues is one method of limiting the amount of work being entered into the system.

FIGURE ON THIS PAGE HAS BEEN DELETED

FIGURE ON THIS PAGE HAS BEEN DELETED

Step 5 - The CPU Time Report and WIN Report can be used to determine those periods of time when TSS and/or WIN programs are using excessive amounts of processor time or, on the other hand, do not appear to be requesting sufficient CPU service.

Step 6 - If the Percent of Memory Time in Queue histogram shows that the average activity is spending more than 30% of its memory time waiting for the processor, this is a strong indication that processing power is a constraining factor. Once again, in order to relieve this constraint, it will be necessary to acquire an additional processor, a faster processor, or else the workload will need to be controlled via the scheduler classes. The CPU Access By SNUMB Report can be used to determine, on an activity-by-activity basis, exactly which activities are being delayed the most due to the lack of processor power.

Step 7 - If the CPU Time Report shows that the percent of system CPU exceeds 30%, this is another indication that the system software is being requested to handle an excessive workload. In all probability, the system queues are increasing to such a size that the system software is expending excessive resources managing its queues.

Step 8 - The CPU Time Report (dispatcher queuing portion) indicates the percentage of time the various system programs spend waiting for the processor and the average queue position of these programs. System programs should not spend more than 2% of the time waiting for service and their average queue position should not exceed 2. If this is not the case, the analyst should ensure that the Urgency Thruput Option is enabled. In addition, the Urgency Report and Excessive Resource Report of the Memory Monitor should be examined to ensure that there is not an excessive number of user activities processing at very high urgencies (see section 14.6.3 for memory evaluation details).

14.6.5 I/O Evaluation. The I/O evaluation will determine whether the mass storage subsystem, or tape channel subsystem is the cause of system degradation. This evaluation requires the user to have processed the Mass Storage Monitor and Channel Monitor data reduction programs.

14.6.5.1 Data Recording. All output from the Mass Store Monitor and Channel Monitor are required. No individual work tables are required, but the user may generate some if he feels that it will help in his analysis.

14.6.5.2 Evaluating the Data. Chapters 7 and 8 provide a fairly detailed description of the procedure to be followed in analyzing the associated reports. In this section, reference will be made to those chapters indicating actual data values that should be used as a reference for comparison.

Step 1 - Read subsection 7.3 and subsections 8.2 and 8.3.

Step 2 - Check the crossbar configuration using the procedure described in subsection 8.2.

Step 3 - Examine the Proportionate Device Utilization Report produced by either the MSM or CM. Check for devices which have significantly higher utilization than other devices in the system. These devices are potential bottlenecks and should be more closely analyzed. It is desirable, even though perhaps not possible, to have equal utilization across all disk packs. Read subsections 7.5.18 and 7.5.20 for further details on this step. Once a pack(s) is identified, further analysis should be performed to determine the actual files being referenced on the pack (see subsection 7.6.1).

Step 4 - The histogram displaying Data Transfer Sizes for TSS Swap Files can give a strong indication of the sizes of TSS subsystems being used. TSS subsystems of over 25K can cause significant increase in overall TSS response, especially if several of these subsystems are being executed simultaneously. If more than 20 percent of the entries in this report fall in the bucket ranges above 25000, this is a strong indication that TSS response might be a problem. This problem can be further confirmed with the CAM.

Step 5 - Seek Elongation - Subsectons 7.5.6 and 7.5.7 describe in detail the reports used to investigate seek elongation problems. An average seek of over 50 cylinders for DSS191s and 100 cylinders for DSU450s should be considered significant.

Step 6 - Analyze the Channel Monitor Idle Report. This report can be generated only if the Idle Monitor was run in conjunction with the Channel Monitor. If the "% of Idle Time During Which I/O Was Active" value exceeds 25%, then substantial benefit may be obtained by eliminating I/O contention. The above value is an indication that even though the CPU is going idle (i.e., has no useful work to perform) there really is potential CPU work available. However, under current conditions, this potential CPU work is being delayed because of I/O contention.

Even though the above figure exceeds 25%, the system may not have sufficient CPU power available to handle the increased work generated by removing the I/O contention. Therefore, the analyst should also check that the "Average System % Idle" figure exceeds 15%. If this proves to be the case, then removal of any I/O contention should prove beneficial. On the other hand, if the figure is lower than 15%, then removal of any I/O contention will probably result in additional CPU contention. The Idle Report will also indicate those devices causing most of the contention. Make a record of the device numbers.

Step 7 - Examine I/O queue length and I/O queue time histograms for individual devices and channels. Queues greater than one and queue times greater than 15 MS should be considered significant. Record those devices with high contention.

Step 8 - If certain devices have been determined as bottlenecks under the procedures described in Steps 1 and 2, the Job Conflict Report should be obtained for those devices following the procedures described in Chapter 8.

Step 9 - Execute the Mass Store Monitor Data Reduction Program. Following the procedure described in subsection 7.6.1 for monitoring a specific device, the analyst should be able to determine the exact files that are causing the contention found under the earlier steps.

Step 10 - Using the CM, it is possible to perform a detailed analysis on channel queuing for a particular job. Details for this procedure can be found in subsections 8.5.13, 8.5.14 and 8.6.10.

Step 11 - This step outlines procedures for relocating files identified as candidates for file relocation. Because of automatic load-leveling activity by the GCOS operating system, an analyst has only limited flexibility for the placement of system, permanent, and temporary files:

a. System Files. The device name on which a system file is to be placed can be specified at system startup. Care should be taken to insure that multiple high-used system files are not placed on the same disk device. If possible, separate high-use system files across disk subsystems. In addition, ensure that SSA Cache memory and FMS cache are enabled to reduce disk I/O activity to certain system files. Details for this analysis can be found in subsections 7.5.9, 7.5.10 and 7.5.19.

b. Permanent Files. The device name for a permanent file can be specified at creation, whether through FMS or the ACCESS subsystem of Time Sharing. Files can be moved by changing their names, creating a new file with the old name, and moving the data. The new file can be created with a DEVICE specification.

c. Temporary Files. The device name for a temporary file can be specified in the second field of the $FILE card in the job control deck. Jobs which run frequently can have their $FILE cards changed. Other jobs can be controlled by policies governing the use of $FILE cards.

Additionally, sites that have different device types may specify preferred device types to be used for temporary files. This procedure will allow activities requiring disk storage to take advantage of higher speed devices.

Step 12 - This step identifies possible seek contention problems attributable to inadequate temporary file space. This procedure uses the SPUTIL feature of the GCOS FMS facility and the Disk Fragmentation Report available at most sites. If such a report is not available, contact CCTC/C751. It is necessary to analyze temporary disk capacity on all disk units rather than just the units identified in previous tuning steps. This analysis is necessary because the disk units exhibiting high activity due to temporary file use often have more available temporary space. The increased utilization of these disk units may be caused by inadequate temporary storage on other disk devices. For this analysis a form as shown in figure 14-5 may prove useful.

a. Report Values. The SPUTIL Report contains the following information on each disk device: (1) device identification, (2) overall capacity, (3) available disk space, (4) disk space dedicated to permanent storage. The Disk Fragmentation Report contains additional information on each disk device: (1) number of disk fragments, (2) average fragment size, (3) maximum fragment size, (4) percentage distribution of fragments by size, and (5) total fragmented space.

b. Form Entry. Enter the device identification for each disk device on the Temporary Storage Test Form. For each disk device enter: (1) the LLINK capacity as indicated by the SPUTIL Report in the Total Capacity column; (2) the temporary capacity from the SPUTIL Report in the Temp Capacity column; (3) the number of fragments from the Disk Fragmentation

DEVICE	TOTAL CAPACITY	TEMP CAPACITY	RATIO 1	RATIO 2	NUMBER OF FRAGMENTS	MAX SIZE	PERCENT FRAGMENTS		TOTAL PERCENT	FRAGMENT
							1-12	13-121		

Figure 14-5. Temporary Storage Test Form

It's gonna swap the whole thing. Say, there's a nice little TSS COBOL subsystem -- only 40K! Time to swap? So soon? Oh yeah, that memory-time quantum. Lot of memory -- not much time. Oh well, a 40K transfer is OK. My swap files are on fast 451s. Friend, you just tied up that device for almost a quarter of a second, and more importantly, a physical channel as well. That's only half of it. TSS swapped it out, you can bet it'll swap it back in. So we're up to .5 sec of physical channel time to free up some TSS user memory. This scenario impacts total system throughput, not just TSS. That's not all the havoc that larger user subsystems create, but that is enough heartburn for now."

14.6.6 Communication Evaluation. The communication evaluation will determine the overall terminal usage of a system. It can also be used to examine the DN355 usage. This evaluation can be done using either the CAM or the GRTS monitor, or both. Figures 14-6 through 14-9 are sample table formats that may be used to display the gathered data.

14.6.6.1 Data Recording. For figure 14-6, the Terminal Session Report is used. All users with TSS subsystems of >35K are recorded. Column 7 is obtained from scanning the Excess Think/Response Time Report. Figure 14-7 is obtained from the Response Time Report. All periods of time when the response for TSS is greater than 15 seconds are recorded. Figure 14-8 is obtained directly from the High Terminal Usage Report. All DAC terminals with over 90 percent usage, except WIN lines, are recorded.

Figure 14-9 is obtained from the H6000-DN355 Reject Report and the Abort Report. The H6000-DN355 Reject Report is used if the average number of reject commands per hour of running time is greater than 50 (total number of reject commands/number of hours in run). Terminals with more than 30 percent of the rejects should also be listed. For the Network Control Program (NCP) disconnects, all NCP 01 line disconnects listed in the Abort Report should be tallied.

14.6.6.2 Evaluating the Data. The following procedure should be followed in order to analyze the data.

Step 1 - TSS response time is dependent upon certain TSS parameters. The TS1 Initial Parameter Report contains current settings of these parameters. The critical parameters are:

a. Initial TS1 Max Size - If operators must increase max size during the day, TSS can slow its processing of user responses until it can grow. This should be set to the normal maximum size TSS reaches.

b. Size Growth/Reduction Factor - These sizes should be identical or growth can be twice reduction.

c. Max Number of Terminals - Should be large enough to satisfy all possible users.

#1	#2	#3	#4	#5	#6	#7
DATE	TERMINAL ID	USERID	LOGON TIME	LOGOFF TIME	TSS SUBSYSTEM SIZE	NUMBER OF EXCESSIVE RESPONSES

Figure 14-6. Large TSS Subsystem Users

DATE	TIME	RESPONSE TIME	IN-RANGE RESPONSE	NUMBER OF OUT-OF-RANGE RESPONSES	NUMBER OF USERS

Figure 14-7. Poor TSS Response Log

DATE	TERMINAL ID	PERCENT USAGE	NUMBER SESSIONS	NUMBER INPUTS	NUMBER OUTPUTS

Figure 14-8. High Usage Terminals

DATE	TERMINAL ID	OPCODE REJECTS	PERCENT REJECTS	NCP DISCONNECTS

Figure 14-9. Terminal/Line Error Report

d. Large Subsystem Size/Wait Time - The average subsystem size should be less than 35K and all large users (greater than 35K) should be penalized.

e. Number Swap Files - On an active TSS system, should be four.

f. Allocated Devices - Swap Files (#S, #T, #U, #V) should be on specific devices, as determined by MSM.

Step 2 - Users of large TSS subsystems cause TSS to swap other users out and therefore generate extra TSS and system I/O overhead. If an excessive number of users are using subsystems of greater than 35K (figure 14-6), these users should be queried as to what subsystem they are using. Any large subsystem that has high usage should be investigated for possible rewrite.

Step 3 - The overall system and TSS response time should be monitored using figure 14-7. Periods of bad response should be checked to ensure the bad responses are not caused by a few (less than 10 percent of all in-range responses) out of range responses. If TSS response is truly poor, a correlation between response and large subsystem usage should be attempted (figure 14-6 and figure 14-7).

Step 4 - Terminals logged onto the system for long periods of time (greater than 75-80 percent of the monitoring session), but having few inputs and/or few inputs/outputs, should be investigated (figure 14-8). Terminals with few inputs, but many outputs, are probably logged onto VIDEO. The number of users on VIDEO should be restricted to one or two per DATANET, due to the buffer load place on the DATANET by VIDEO. If more users are required, monitor output terminals should be used instead of more VIDEO users. Terminals with few inputs and few outputs may be being used just to keep a terminal logged on. This practice causes unnecessary TSS size and processing overhead.

Step 5 - Terminals with a large number of Opcode Rejects (figure 14-9) are an indication of possible line problems. The terminals and lines should be checked for noise and transmission errors. Numerous NCP disconnects per day (figure 14-9) is usually an indication of line or IMP problems. NCP should have fewer than five disconnects per day.